UNIVERSITY OF KANSAS | DIGITAL STRATEGY | ADAPTIVE INTELLIGENT ECOSYSTEM

# admissions.ku.edu/explore — Block v2

## Optimization & Style-Update Analysis

A semantic and typographic refactor of the /explore dwell-driver block: three-family type, article-level entities, and a single linked JSON-LD graph — measured across all eight of our categories.

GEO | SEO | AEO | UX | WCAG 2.1 AA | Conversions | Accessibility | UI

**What this deck covers**

1. The five v2 changes and how each one serves entity extraction
2. Apples-to-apples: the original v1 block vs. the v2 refactor, across all eight categories
3. The schema, honestly — what each type produces, and why the linked graph matters
4. Exactly where we monitor the effect in GA4, Search Console, Ahrefs, and Microsoft Clarity
5. Our measurement timeline, 90-day KPIs, and a full-code reference appendix

Single-paste KUTE deliverable: ku-explore-v2.html | Replaces #ku-explore-v1

# What v2 Changes — Refactor Summary

Same content, same answer copy; five structural upgrades

### 1. TYPOGRAPHY — THREE FAMILIES

United Sans is demoted to eyebrows, chips, CTA buttons, tags, and small labels only. The H2, section H3s, card titles, and stat numbers move to the Roc Grotesk / Halyard Display head family in sentence case. A runtime IIFE reads the live page H1 and syncs the head font and weight so our block always matches the theme.

### 2. SEMANTICS — ARTICLE ENTITIES

Feature cards and campus districts become <article> elements — self-contained entities an AI extractor can lift cleanly. Heading wrappers are removed, buttons are direct children, and the text-to-code ratio rises.

### 3. RHYTHM — PY-5 CADENCE

Outer padding moves to a clamp(3rem, 5vw, 4rem) rhythm with a top and bottom hairline rule, so the block reads as part of the native KUTE paragraph flow rather than an inserted island.

### 4. HARDENING — CKEDITOR-SAFE

Buttons are fully reset (appearance:none), every DOM write is null-guarded, panel copy lives in the initial HTML so AI crawlers read it, and list-style plus sr-only rules are scoped inside the block.

### 5. SCHEMA — ONE LINKED GRAPH

The four types now share one @graph, cross-referenced by @id: WebPage.about points to the KU organization, .mainEntity to the FAQPage, and .breadcrumb to the BreadcrumbList. dateModified is added and every answer stays in parity with the visible copy. This is the connective tissue AI extraction relies on.

**Type tokens after the refactor**

| Token | Family | Now used for |
|---|---|---|
| --kx-f-ey | United Sans | Eyebrows, chips, CTAs, tags, labels |
| --kx-f-head | Roc Grotesk / Halyard Display | H2, section H3s, card titles, stat numbers |
| --kx-f-body | Halyard Text | Body copy, ledes, descriptions |

# Apples-to-Apples: v1 vs. v2

A near like-for-like read of what changed for a visitor, a crawler, and an AI engine

| Dimension | Original #ku-explore-v1 | Refactored #ku-explore-v2 |
|---|---|---|
| Feature markup | Cards and districts as <ul>/<li> list items | Self-contained <article> entities for clean AI extraction |
| Headings | <h3> nested inside  wrappers | <h3> direct, id and scroll-margin on the heading itself |
| Typography | United Sans on H2, every H3, card titles, and stat numbers | United Sans demoted to eyebrows/labels/CTAs; head family on H2, H3, titles, stat numbers |
| Theme sync | None — fixed font stack | Runtime IIFE syncs head font + weight to the live page H1 |
| Structured data | Four correct but disconnected @graph nodes | Five linked nodes; WebPage about / mainEntity / breadcrumb resolve by @id; dateModified added |
| Emphasis | About 17  instances | 10  — lighter, more scannable answer-first copy |
| Buttons | Wrapped in <p>; theme-dependent styling | Fully reset (appearance:none), CKEditor-safe, injected-<p> defense |
| Instrumentation | Toggle behavior only | Single explore_interaction dataLayer event on disclosures, FAQ, and chips |

# Strategic Alignment Matrix

How the refactor maps to each of our eight optimization categories, and the signal we watch

| Category | What v2 does | Signal we watch |
|---|---|---|
| GEO | Linked @graph ties the page to the KU entity, its FAQ, and its breadcrumb; panel copy sits in initial HTML for JS-free AI crawlers | Clarity Citations; GA4 AI Assistant sessions |
| SEO | Semantic article/H2/H3 outline, high text-to-code ratio, valid linked structured data | GSC impressions on question queries; breadcrumb appearance |
| AEO | Six-question FAQPage mirrors visible Q&A; speakable targets title, lede, answers | Featured-snippet / Position Zero ownership in Ahrefs |
| UX | Mobile-first grids, quick-nav chips, disclosure panels cut scroll friction | Clarity scroll depth and chip/panel use |
| WCAG 2.1 AA | aria-expanded/controls, keyboard-native buttons, 3px lake focus ring, image dimensions prevent layout shift | Manual audit plus DubBot |
| Conversions | Verb-first CTAs in KU Fire-and-Ice buttons at each entity | GA4 CTA click events |
| Accessibility | Reduced-motion override on every animation, decorative icons aria-hidden, scoped sr-only rule | Assistive-technology pass |
| UI | KU Blue/Crimson/Lake tokens with hard-coded fallbacks, gradient Fire-and-Ice stat underline | Brand review sign-off |

# The Schema, Honestly

Four types, three jobs — set expectations so we watch the right signal for each

| Schema type | What it does for us | Honest status in 2026 |
|---|---|---|
| BreadcrumbList | Clarifies the Admissions > Explore hierarchy | Still produces a standard breadcrumb rich result — the one visible SERP payoff here |
| FAQPage | Maps our answers to real questions for AI Overviews and assistants | Valid and parsed, but FAQ rich results were removed May 7, 2026 — no snippet, no Enhancements report |
| CollegeOrUniversity | Strengthens how engines and AI recognize KU as an entity | Entity signal for knowledge panels and AI answers; never a standalone rich result |
| WebPage + speakable | Names the page and flags voice-friendly passages | Comprehension and voice aid; no dedicated report |

### Why the linked graph is the real upgrade

v1 shipped four correct but floating nodes, so an engine had to infer that the FAQ, the breadcrumb, and the KU entity all belonged to this page. v2 states it outright through @id references. That is exactly the connective tissue AI extraction and knowledge-graph reconciliation rely on — and it costs nothing in visible content, since our answer copy, links, and images are unchanged. We treat schema as low-cost reinforcement of strong visible content, not a magic switch.

# Research Foundations & Citations

Every decision traces to a primary source below

- **Answer-first content wins AI answers.**
  The first peer-reviewed GEO study (Princeton / Georgia Tech / IIT Delhi, KDD 2024) found that citations, statistics, and quotations lift visibility in generative answers by roughly 30-40%. Our FAQ and stat blocks put concrete figures in extractable positions. [1]

- **AI crawlers read HTML, not JavaScript.**
  Vercel / MERJ analysis of 500M+ fetches found GPTBot, ClaudeBot, and PerplexityBot execute zero JavaScript. v2 keeps all panel copy in the initial HTML, so nothing an AI needs is trapped in a script. [2]

- **FAQ rich results are gone; the markup still works upstream.**
  Google removed FAQ rich results on May 7, 2026, and the Search Console FAQ report is sunsetting. Both types remain valid and parsed; the visible Q&A does the load-bearing work. [3]

- **Breadcrumb and speakable still pay off.**
  Breadcrumb schema continues to produce a standard rich result, and speakable supports voice answers — the two still-live wins in our graph. [3]

- **Entity clarity drives AI recognition.**
  Google's guidance: no special schema is required for AI Overviews, but structured data should match visible content. Our linked CollegeOrUniversity node reinforces KU as a known entity. [4]

- **Structured data must mirror the page.**
  Schema.org linked-graph modeling (@id references) is the documented way to connect a WebPage to its organization, FAQ, and breadcrumb. [5]

**Sources**


[1] Aggarwal et al., Generative Engine Optimization, ACM KDD 2024 — arxiv.org/abs/2311.09735
[2] Vercel & MERJ, The Rise of the AI Crawler — vercel.com/blog/the-rise-of-the-ai-crawler
[3] Google Search Central, FAQ rich-result changes (2023-2026) — developers.google.com/search/blog
[4] Google, structured data for AI features guidance (2026); [5] Schema.org linked-data modeling (@id).

# Where We Monitor: Google Search Console

Confirm the schema, then watch the one type that still draws a feature

### A. VALIDATE THE MARKUP

URL Inspection on /explore confirms all four types are detected and valid. The Rich Results Test validates the BreadcrumbList (a live feature) — no FAQ or entity Enhancements report is expected, by design.

### B. WATCH QUERIES & APPEARANCE

Performance > filter to the URL. Track our six FAQ questions as queries and branded terms for the entity signal. Use the Search Appearance filter to confirm the breadcrumb trail is rendering in results.

### C. THE NEW AI VIEW

Google launched a dedicated Generative AI performance report on June 3, 2026, under Performance > Search results in its own tab, showing how often our URLs appear in AI Overviews and AI Mode by page, country, device, and date.

### Read it correctly

- Impressions only — no clicks, CTR, or query breakdown yet.
- A breakout of existing data: AI impressions were always in our totals, so aggregate numbers do not change.
- Phased rollout starting with UK properties; data begins about May 18, 2026 with no backfill — so baseline the day it appears.
- A flat organic line with rising AI impressions means our content is being picked up even as clicks shift.

Sources: developers.google.com/search/blog/2026/06/gen-ai-performance-reports; support.google.com/webmasters/answer/16984139 (Generative AI performance report, June 3, 2026).

# Where We Monitor: Google Analytics 4

Engagement first, then AI-referred sessions — with one setup step that closes a gap

### A. ENGAGEMENT ON /EXPLORE (PRIMARY KPI)

Reports > Engagement > Pages and screens, filter to /explore. Watch Average engagement time per active user against our pre-launch baseline; the chips, disclosure panels, and animated stats should move this first. Register explore_interaction (fired on .kx-disc-btn / .kx-faq-btn / .kx-chip via GTM-K75HJX66) and watch Event count per session.

### B. NATIVE AI ASSISTANT CHANNEL

GA4 added a native AI Assistant channel on May 13, 2026 that auto-classifies sources like ChatGPT, Gemini, and Claude with no setup, tagging sessions medium = ai-assistant. Find it in Reports > Acquisition > Traffic acquisition.

### C. CLOSE THE GAP

The native channel is forward-only (no backfill), Perplexity and Copilot still land in Referral, and AI Overviews clicks count as Organic Search. Keep a custom channel group with an AI-source regex ordered ABOVE Referral — it applies retroactively and catches the sources the native channel misses.

**Custom channel group — source regex**

```
chatgpt\.com|openai\.com|perplexity\.ai|
claude\.ai|gemini\.google\.com|
copilot\.microsoft\.com
```

Place this rule above Referral in the channel-group order. Below Referral it never fires, because Referral's broad rule matches AI domains first. 20-40% of AI visits arrive with no referrer and fall into Direct — treat the channel as a floor, not a census.

Sources: digitalapplied.com/blog/ga4-ai-assistant-channel-2026 (native channel, May 13, 2026); gaoptimizer.com / seosherpa.com (keep a custom channel group; Perplexity gap; referrer-less Direct).

# Where We Monitor: Ahrefs & Microsoft Clarity

Cross-platform AI visibility, and first-party proof the page is being cited

## Ahrefs — cross-platform AI visibility

- Rank Tracker: add our explore/campus/location queries; the AI Overview SERP-feature filter flags when a keyword triggers an AI answer and whether we appear (base plans, no add-on).
- Brand Radar: tracks Mentions, Citations, and AI Share of Voice across six indexes — AI Overviews, AI Mode, ChatGPT, Perplexity, Gemini, and Copilot. Enter 'University of Kansas' and our FAQ questions as custom prompts.
- Site Explorer > AI citations module: the pages winning citations over time. Expect movement on a 60-90 day horizon.

Sources: ahrefs.com/blog/how-to-track-ai-overviews; help.ahrefs.com (Brand Radar: six indexes, mentions/citations/AI Share of Voice).

## Microsoft Clarity — behavior + citations

- AIPlatform / PaidAIPlatform channel groups (Channels card) isolate AI-referred sessions; filter recordings and heatmaps on /explore to that segment.
- Confirm visitors reach the stats and FAQ; watch dead / rage clicks on the widgets as our CKEditor-regression alarm.
- The Citations dashboard went generally available May 13, 2026: grounding queries, cited pages, share of authority, and AI-referred traffic — first-party proof the page is pulled into AI answers.
- Scope: Microsoft surfaces (Copilot, Bing) only. Verify the domain by connecting Google Search Console or Bing Webmaster Tools.

Sources: learn.microsoft.com/clarity/ai-visibility/ai-citations; clarity.microsoft.com/blog (Citations GA, May 13, 2026).

# Our 90-Day Measurement Plan

One baseline this week, then a fixed cadence — expecting the usual 6-8 week lag before AI-surface movement

| Metric | Tool & location | Baseline | 90-day target | Cadence |
|---|---|---|---|---|
| Avg engagement time, /explore | GA4 > Engagement > Pages and screens | Capture now | Upward vs. baseline | Weekly |
| explore_interaction / session | GA4 events (GTM, both containers) | 0 (new) | >= 0.35 | Weekly |
| Scroll depth to stats/FAQ | Clarity scroll heatmap, /explore | Capture now | >= 50% of sessions | Weekly |
| Breadcrumb appearance | GSC > Search Appearance | Confirm at publish | Rendering in results | Biweekly |
| Generative AI impressions | GSC > Performance > Generative AI tab | Set when tab arrives | Upward trend | Biweekly |
| AI Assistant + AI-regex sessions | GA4 channel groups | Set at publish | Upward trend | Biweekly |
| FAQ-question snippet ownership | GSC Performance + Ahrefs | Unranked | Own 2 of 6 questions | Monthly |
| AI citations for the page | Ahrefs Brand Radar + Clarity Citations | 0 tracked | Cited on >= 2 platforms | Monthly |
| Branded search demand | GSC branded queries; Brand Radar | Current level | Positive lift | Monthly |

How we read it: AI answers are non-deterministic, so we track ranges and trends, not single-day values. Engagement moves within 2-3 weeks of publish (user-behavior driven); schema-dependent AI-surface metrics follow the 6-8 week indexing lag we have documented on prior KU deployments. Rising crawl activity or Clarity citations while referrals stay flat is visibility-without-clicks — expected for informational pages, and exactly the durable AI-surface footprint the Adaptive Intelligent Ecosystem targets. All of this rolls up into our three reporting pillars: Discoverability, Engagement, and Attribution.

# Appendix — Full v2 Source (Reference Copy)

Complete single-paste deliverable: ku-explore-v2.html. Zoom to read; the .html file is canonical.