UNIVERSITY OF KANSAS | DIGITAL STRATEGY | ADAPTIVE INTELLIGENT ECOSYSTEM

online.ku.edu/academics — Block v3

Optimization Analysis & Measurement Playbook

Responds to design feedback of 6.29.26: py-5 section rhythm, three-family typography, reduced bolding, and connected-panel responsive cards — plus new GEO hardening that puts 100% of copy in the initial HTML.

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

What this deck covers

1. The four 6.29.26 design-feedback items and how v3 implements each one
2. Apples-to-apples: the original static Drupal block vs. the v3 adaptive block, across all eight of our categories
3. The research and sources behind every decision, with citations
4. Exactly where we monitor schema and engagement impact in GA4, Search Console, Ahrefs, and Microsoft Clarity
5. Our 90-day KPI plan and a full-code reference appendix

Prepared July 6, 2026 | Single-paste KUTE deliverable: ku-academics-v3.html | Replaces #ku-acad-v2

What v3 Changes — Design Feedback of 6.29.26, Implemented

Same content, same schema parity; four structural design revisions plus one new layer of GEO hardening

1. PADDING — PY-5 RHYTHM

Outer block padding now runs clamp(3rem, 5vw, 4rem) to match Drupal's native py-5 (3rem) cadence, and inter-section gaps widen to clamp(3rem, 6vw, 4.5rem). A top and bottom hairline rule on the inner wrap mirrors the visual cadence of the native KUTE paragraphs above and below the block, so v3 reads as part of the page rather than an insert.

2. FONT DIVERSITY — THREE FAMILIES

United Sans (condensed) is demoted to eyebrows, buttons, and small labels only. H2s, reason and step titles, and stat numbers move to the Roc Grotesk / Halyard Display family in sentence case. A runtime IIFE reads the computed font-family and font-weight from the live page H1 (h1.page-title--halyard) and writes them to --a-f-head / --a-f-head-w, so headings always match the KU Sunflower theme exactly — our confirmed cross-block pattern.

Token	Family	Used for
--a-f-display	United Sans	Eyebrows, buttons, small labels
--a-f-head	Roc Grotesk / Halyard Display	H2s, panel & step titles, stat numbers
--a-f-body	Halyard Text	Body copy, ledes, descriptions

3. BOLDING — REDUCED ~20%

In-copy emphasis drops from 14 to 11 instances (a 21% reduction), easing visual and screen-reader load while keeping the most scannable facts: faculty parity, AAU and Carnegie R1 status, the 4/8/16-week formats, and the any-number-of-credits transfer message.

4. CARDS — CONNECTED DESKTOP PANELS

Mobile keeps the card treatment (borders, shadows, radius) that tests well on small screens. At desktop widths the same markup dissolves into a continuous connected panel: shared hairline dividers, no per-item shadow, one container border and shadow. Applied to the Why-KU reason list at 820px and the three degree-completion steps at 880px, with a crimson active edge preserving the Fire-and-Ice hover logic. Implemented mobile-first, cards as the base and min-width media queries as the override.

NEW IN V3: GEO HARDENING

In v2 the Synchronous and Collaborative format descriptions existed only inside a JavaScript object — invisible to AI crawlers, which fetch raw HTML and do not execute scripts. v3 renders all three format panels in the initial HTML and JS only toggles the hidden attribute. Tablist and format panels also normalize their hidden state on init, self-repairing if CKEditor strips attributes on save.

Apples-to-Apples: Original Static Drupal Block vs. v3

Near like-for-like comparison of what the CMS ships to a visitor, a crawler, and an AI engine

Dimension	Original built-in Drupal block (static cards)	v3 adaptive block
Markup	Utility-class div nesting; meaning carried by visual styling only; low text-to-code ratio	Semantic section/h2/p structure, scoped classes, flat CKEditor-tolerant markup, high text-to-code ratio
Structured data	None	@graph: CollegeOrUniversity (AAU, sameAs, Wikidata) + WebPage (speakable, dateModified) + 5-question FAQPage kept in copy parity
AI-crawler visibility	Copy visible, but no entity or Q&A scaffolding for extraction	100% of copy in initial HTML (v2's JS-only format text moved into markup); direct-answer lede in first 60 words; FAQ answers mirror on-page copy
Dwell-time drivers	None — read-only cards (page baseline: 18s avg engagement)	Three interactive widgets: Why-KU tablist, Build-Your-Week format toggle, credit estimator with live progress bar
Accessibility	No keyboard model, no reduced-motion handling, contrast unaudited	Full ARIA tablist with arrow keys, aria-live outputs, 41.6px targets (WCAG 2.2 min is 24px), prefers-reduced-motion override, no-JS fallback shows everything
Performance	Theme-dependent; images/icons as assets	Zero additional HTTP requests; CSS-mask vector icons; IntersectionObserver-gated transform/opacity animation (compositor-friendly, INP-safe); no frameworks
Instrumentation	Pageviews only	Click surfaces (.acj-tab / .acj-fmt / .acj-cr) ready for an academics_interaction GTM event in both containers (GTM-K75HJX66, GTM-MG3V335)

Strategic Alignment Matrix — All Eight Optimization Categories

Every v3 decision maps to at least one category and one observable signal in our tools

Category	What v3 does	Signal we watch
GEO	Entity graph (KU as AAU/R1 org), 100% server-delivered copy, statistics and named entities in-copy per the Princeton GEO findings	GSC Generative AI impressions; Ahrefs Brand Radar citations; Clarity AI Citations
SEO	Semantic headings, direct-answer lede, dateModified freshness signal, internal link to /undergraduate-programs	GSC Performance queries & position for /academics
AEO	5-question FAQPage in strict copy parity; speakable spec on the answer block; question-format H2s	Featured-snippet ownership in Ahrefs; People-Also-Ask coverage
UX	Connected desktop panels reduce visual noise; mobile-first card layout; interactive widgets replace passive scanning	GA4 avg engagement time 18s toward the 44-52s norm; Clarity scroll depth
WCAG 2.1 AA	Contrast-checked tokens, focus-visible outlines, reduced-motion override, keyboard-complete tablist	DubBot scans; manual keyboard/AT audit
Conversions	Verb-first estimator note routes to program finder; removed redundant stacked CTA so the page's finder stays primary	GA4 key events; Slate RFI attribution on gradapply
Accessibility	aria-live announcements, 41.6px targets, sr-safe labels, no-JS content parity for assistive tech and text browsers	Clarity dead/rage clicks on widgets (also our CKEditor regression alarm)
UI	KU brand tokens, Fire-and-Ice hover states, three-family type system synced to the live theme H1	Design review sign-off; Clarity recordings

Research Foundations & Citations

The external evidence behind each decision — every claim traceable to a primary source below

- **GEO tactics [1]** The first peer-reviewed GEO study (Princeton/Georgia Tech/IIT Delhi, KDD 2024; 10,000 queries) found that adding statistics, citations, and quotations lifts visibility in generative-engine answers by roughly 30-40%. v3 keeps concrete figures (AAU ~70 institutions, 4/8/16-week formats, 120-credit math) in extractable positions.
- **AI crawlers & JS [2]** Vercel/MERJ analysis of 500M+ GPTBot fetches found zero JavaScript execution; GPTBot, ClaudeBot, and PerplexityBot read raw HTML only. This is why v3 moves the Synchronous and Collaborative copy out of the JS object and into the initial HTML — the single biggest GEO fix in this revision.
- **FAQ schema status [3]** Google limited FAQ rich results to authoritative government and health sites in August 2023, so we do not expect an Enhancements report — the FAQPage markup instead works upstream, helping search and answer engines parse our Q&A structure for AI Overviews, AI Mode, and assistants. Validation happens in the Rich Results Test / Schema.org validator, not a GSC report.
- **WCAG 2.2 targets [4]** WCAG 2.2 (W3C Recommendation) adds Target Size (Minimum) at 24x24 CSS px; our stepper buttons render at 41.6px and tabs larger still, exceeding the AA bar while we formally target 2.1 AA. Focus-visible outlines are preserved inside the new connected panels via negative outline-offset.
- **Core Web Vitals [5]** Interaction to Next Paint (INP) replaced FID as a Core Web Vital in March 2024. v3 stays vanilla-JS, event-delegated, and animates only transform/opacity behind IntersectionObserver, keeping main-thread work minimal on a page already running two GTM containers.
- **Copy grounding [6]** RNL's 2024 online student research shows fewer than 10% of online bachelor's students arrive with zero credit and only about 21% with 60+, grounding the any-number-of-credits message and the estimator's default of 60 earned credits.

Sources

[1] Aggarwal et al., GEO: 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, Changes to HowTo and FAQ rich results (Aug 2023) — developers.google.com/search/blog/2023/08/howto-faq-changes
[4] W3C, WCAG 2.2, SC 2.5.8 Target Size (Minimum) — w3.org/TR/WCAG22/
[5] Google web.dev, Interaction to Next Paint — web.dev/articles/inp
[6] RNL, 2024 Online Student Recruitment Report (project source library)

Where We Monitor: Google Search Console

Two workflows — the brand-new Generative AI report for schema/AI impact, and the classic Performance report for organic lift

A. GENERATIVE AI PERFORMANCE REPORT (NEW, JUNE 3 2026)

Google now ships a dedicated report showing how often our URLs appear inside AI Overviews and AI Mode (plus a sibling report for Discover). It lives at Performance > Search results > Generative AI tab once our property is included.

- Filter Page = /academics and chart impressions weekly; this is the most direct read on whether the schema and answer-first copy are earning AI-surface visibility.
- Impressions only for now — no clicks, CTR, or queries yet; Google says more metrics are coming. Pair with GA4 for the click side.
- Rollout is phased (UK properties first); until it appears for us, AI impressions remain blended into our normal Performance totals, so those totals will not change when the tab arrives.
- Baseline the number the day it appears — every future metric gets measured against that baseline.

B. CLASSIC PERFORMANCE REPORT

- Performance > Search results, filter Page = /academics; watch impressions and average position for our five FAQ question queries (e.g., 'how do online classes work at KU', 'transfer credits to KU online').
- Compare 28-day windows before/after publish; expect the 6-8 week indexing lag we have seen on prior schema deployments before judging impact.
- Watch for the widening impressions-vs-clicks gap on question queries — the standard signature of AI Overviews answering above our listing.
- URL Inspection after publish confirms Google sees the FAQPage and Organization markup in the rendered HTML; the Rich Results Test / Schema.org validator is our pass/fail check, since FAQ no longer generates an Enhancements report for .edu sites.
- Watch branded queries ('KU online', 'University of Kansas online degree') for lift — our leading indicator that AI mentions are driving brand searches.

Sources

Google Search Central Blog, Introducing Search Generative AI performance reports (June 3, 2026) — developers.google.com/search/blog/2026/06/gen-ai-performance-reports
Google Search Console Help, Generative AI performance report — support.google.com/webmasters/answer/16984139

Where We Monitor: Google Analytics 4

Three views — the new AI Assistant channel, a custom channel group to close its gaps, and page-level engagement

A. AI ASSISTANT CHANNEL (NATIVE, MAY 13 2026)

- Reports > Acquisition > Traffic acquisition, primary dimension = Session default channel group. Sessions referred by ChatGPT, Gemini, and Claude now auto-classify as 'AI Assistant' (medium: ai-assistant).
- Known gaps: Perplexity and Copilot still land in Referral; AI Overviews clicks count as Organic Search; there is no backfill before the launch date.
- Dark-traffic caveat: a large share of AI referrals arrive with stripped referrers (mobile apps, ChatGPT Atlas) and show as Direct — treat the channel as a conservative floor, not a census.

B. CUSTOM 'AI REFERRALS' CHANNEL GROUP

Admin > Data display > Channel groups > create 'AI Referrals' with Source matches regex:

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

Drag the rule ABOVE Referral (GA4 evaluates top-down; below Referral it never fires). Applies retroactively in reporting, closing the Perplexity/Copilot gap.

Sources

Google Analytics: AI Assistant default channel (May 13, 2026) — coverage: digitalapplied.com/blog/ga4-ai-assistant-channel-2026-measure-ai-traffic-playbook
Custom channel group setup and rule ordering — discoveredlabs.com/blog/how-to-track-chatgpt-perplexity-and-ai-overviews-traffic-in-ga4-without-guessing

C. ENGAGEMENT ON /ACADEMICS (OUR PRIMARY KPI)

- Reports > Engagement > Pages and screens, filter Page path = /academics, column: Average engagement time per active user. Baseline 18s; target the 44-52s norm. Review weekly with an annotation on publish date.
- Register the academics_interaction GTM event (clicks on .acj-tab / .acj-fmt / .acj-cr) and watch Event count per session on this page — this separates 'widgets are being used' from raw time.
- In Explore, build a free-form report: dimension Landing page + Session default channel group, metrics Sessions, Engaged sessions, Average engagement time — this shows whether AI-referred visitors to /academics out-engage organic ones.
- Optionally mark academics_interaction as a key event to make it visible in standard Acquisition reporting alongside RFI conversions.

Where We Monitor: Ahrefs + Microsoft Clarity

Ahrefs covers upstream AI/SERP visibility across engines; Clarity covers on-page behavior plus first-party Microsoft citation data

AHREFS

- Site Explorer > enter online.ku.edu > AI citations module: our confirmed AI Overview citations and their trend (tick the purple AI Overview checkbox on the Overview history chart).
- Rank Tracker: filter tracked keywords by SERP feature = AI Overview to see which of our target queries trigger AIO and whether we appear; add our five FAQ questions as tracked keywords.
- Brand Radar (six indexes: AI Overviews, AI Mode, ChatGPT, Perplexity, Gemini, Copilot): benchmark KU's mentions, citations, and AI Share of Voice against peer institutions; add our five FAQ questions as custom prompts and check monthly.
- Organic keywords > 'owns snippet' filter tracks featured-snippet capture for the question-format H2s — our Position Zero scoreboard.

MICROSOFT CLARITY

- Dashboard > Traffic card > Channels tab: the AIPlatform channel isolates sessions arriving from ChatGPT, Claude, Gemini, Copilot, and Perplexity — then apply it as a segment to heatmaps and recordings.
- AI Citations dashboard (GA since May 2026; Dashboards > AI Visibility > Citations): first-party Microsoft retrieval data showing grounding queries, cited pages, share of authority, and AI-referred traffic for our verified domain. Connect GSC once to verify. Microsoft ecosystem only (Copilot/Bing/ChatGPT-via-Bing).
- Filter to /academics: scroll-depth heatmaps confirm visitors reach the credit estimator; click heatmaps confirm the tabs, format pills, and steppers are actually used — our dwell-time drivers working as designed.
- Dead-click and rage-click alerts on the widgets double as our CKEditor-regression alarm: if a save strips the JS hooks, frustrated clicks appear here within days.

Sources

Ahrefs, How to Track AI Overviews — ahrefs.com/blog/how-to-track-ai-overviews | Ahrefs Help, What is Brand Radar — help.ahrefs.com/en/articles/11064852
Microsoft Learn, AIPlatform channel — learn.microsoft.com/en-us/clarity/insights/ai-channel-group | Clarity Blog, Citations Now Generally Available — clarity.microsoft.com/blog

Our 90-Day Measurement Plan

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

Metric	Tool & location	Baseline	90-day target	Cadence
Avg engagement time, /academics	GA4 > Engagement > Pages and screens	18s	44-52s norm	Weekly
academics_interaction events / session	GA4 events (GTM, both containers)	0 (new)	>= 0.35	Weekly
Scroll depth reaching estimator	Clarity scroll heatmap, /academics	TBD at publish	>= 50% of sessions	Weekly
Generative AI impressions, /academics	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 rankings & snippet ownership	GSC Performance + Ahrefs 'owns snippet'	Unranked	Own 2 of 5 snippets	Monthly
AI citations for 5 FAQ prompts	Ahrefs Brand Radar + Clarity Citations	0 tracked	Cited on >= 2 platforms	Monthly
Branded search demand	GSC branded queries; Brand Radar Search Demand	Current level	Positive lift	Monthly

How we read it: AI answers are non-deterministic, so we track ranges and trends rather than single-day values. Engagement metrics should move within 2-3 weeks of publish (they are user-behavior driven); schema-dependent AI-surface metrics follow the 6-8 week lag we have documented on prior KU deployments. If crawl activity or Clarity citations rise while referrals stay flat, that is visibility-without-clicks — expected for informational pages, and exactly the durable AI-surface footprint the Adaptive Intelligent Ecosystem strategy targets.

Appendix — Full v3 Source (Reference Copy)

Complete single-paste deliverable: ku-academics-v3.html. Zoom to read; the .html file is the canonical copy for deployment.









