Reflections on the Adoption and Integration of AI Tools in Academic Libraries

Amber Wu
Michelle S. Navarro
California State Polytechnic University Pomona

Abstract

As universities expand artificial intelligence (AI) initiatives and library vendors add AI-enabled features to research platforms, academic libraries are increasingly asked to evaluate, explain, troubleshoot, and support tools they may not have selected or fully shaped. AI integration into databases raises important questions about how library workers can evaluate and support these tools with care, accountability, and professional judgment when vendor documentation, campus governance, pricing models, and long-term support structures are still developing.

This conceptual paper examines vendor-embedded AI in academic library contexts from the perspective of early-career librarians and argues that academic libraries are well positioned to assess vendor-embedded AI because library workers already evaluate resources intended to support their campus communities. However, libraries risk being placed in passive roles when vendors or campus initiatives define the timeline, purpose, and value of AI tools before library-centered evaluation has occurred. Rather than positioning academic libraries as either enthusiastic adopters or resisters of AI, this paper encourages a critically engaged middle position.

Article Type: Conceptual Paper

Introduction

Artificial intelligence (AI) is no longer a concept the academic community encounters by choice. Higher education often discusses AI at a broad or aspirational level through strategic planning, workforce preparation, and classroom experimentation. As a result, students, faculty, administrators, and staff increasingly use AI services or tools through their everyday work, institutional conversations, required training, and new policies with expectations to use or support emerging tools. However, if a university acquires an emerging AI product for campus-wide use, providing access is only one part of implementation. What may begin as a university strategic priority or a common tool or feature for users becomes implementation and maintenance work for the employees expected to put it into practice. Information technology staff set up and manage technical infrastructure, security, access, and system integration, but further successful implementation still requires respective collaboration across campus when it comes to research needs and whether its use meets academic standards. Within that shared implementation structure, academic libraries then take on responsibilities that connect the technology to research, instruction, and user support. Library workers will feel the pressure, if not already, to evaluate research-related features, explain the tool to users, incorporate it into instruction and services, troubleshoot

Endnotes: The Journal of the New Members Round Table
Volume 14, Issue 1, 2026
pp. 63-71
Published by the American Library Association

ISSN: 2159-0591

information-related concerns, and assess how it aligns with institutional and professional expectations. Ultimately, academic libraries inherit responsibility even when they were not involved in selecting the tool or establishing the broader campus strategy. But because campus-wide AI adoption requires shared responsibility across the institution, we argue that academic libraries should not settle for forms of participation that begin only after vendors or campus leaders have already defined a tool's purpose, value, timeline, and support expectations. Rubrics, demonstrations, and short trials can help libraries evaluate a product, but they remain reactive when libraries cannot question its underlying assumptions, request changes, clarify responsibilities, or influence whether and how it will be adopted. Drawing on reflections about vendor evaluation, market sustainability, open-source artificial intelligence, and staff capacity, this paper examines how libraries can move from reviewing finished products to shaping the conditions under which artificial intelligence tools are selected, implemented, and sustained.

Academic Libraries within Institutional AI Adoption

As the global expansion of AI technologies accelerated after the public release of ChatGPT in late 2022, higher education institutions across the United States experienced growing expectations to create or expand AI-related courses, certifications, and curricular initiatives. Recent national analyses indicate that completions of AI degrees and certificates grew nearly tenfold between 2016 and 2023, alongside the rapid expansion of generative AI course enrollments and non-degree credential programs (Validated Insights Inc., 2025). This shift reflects an industry growth trajectory driven by labor market demand for AI skills and competencies. Alongside this curricular expansion, institutions moved quickly to establish institutional policies and operational responses to AI adoption. One analysis of U.S. universities found that many institutions had already developed formal institutional responses to AI between late 2022 and 2023, demonstrating how quickly campuses began responding to the larger AI market and its pressures (McDonald et al., 2025).

In the California State University (CSU) system, the momentum of incorporating AI tools on campus became visible through a system-wide initiative announced in 2025. The initiative promised hands-on experience and opportunity to use AI tools, training, and workforce development opportunities across the system and frames AI integration as part of a broader institutional priority (California State University, 2025). The CSU is an example that shows how quickly institutional momentum can move from broad strategic language into local implementation work. Yet as institutions race to adopt and integrate AI, questions of infrastructure, documentation, evaluation, and long-term readiness remain underexamined. For example, a new artificial intelligence feature may appear in a database, a discovery platform may introduce a research assistant, or a vendor demonstration may present a tool as the next stage of search. It is within this gap that academic libraries find themselves both positioned and pressured to respond.

Libraries should avoid being pulled into passive or reactive AI adoption. Libraries should not simply be on the receiving end of AI tools after vendors or campuses have already shaped the purpose, timeline, and support expectations. Instead, libraries should be part of the conversation earlier so they can ask questions, name concerns, define expectations, and help determine whether a tool is worth the investment. Libraries should also be cautious about overinvesting in AI products before direction, documentation, value, and long-term sustainability are clear. The discussion of open-source AI offers an inspiration point because it shows why definitions, transparency, and shared principles matter before tools become normalized. Together, the examples in this paper point toward a middle position that allows academic libraries to question AI tools carefully without refusing the possibility that some tools may become useful.

Another AI Rubric is Too Passive to be a "Solution"

Rubrics are popular in libraries for many reasons, and as the AI field has flooded the library's technology market, it is natural that critically minded librarians have turned to rubrics to evaluate AI products. Academic libraries are often positioned to assess these tools because librarians already have expertise in acquisitions, information literacy, research support, and the critical evaluation of information systems (Nogueira et al., 2025). Rubrics allow reviewers to test new tools quickly and consistently, clarify thinking about data privacy, bias, accessibility, and impact, and communicate the reasoning behind decisions to administrators and governing bodies. These tools appear in a wide range of contexts, including institutional library guides, instructional initiatives, professional organizations, and industry-facing evaluation models (Caico et al., 2024; Mackie & Aspenlieder, 2024).

For example, Caico et al. (2024) provide questions drawn from the ACRL Framework to help librarians decide whether an AI tool is appropriate for their work. Their questions emphasize context, use, ethics, authority, and reflection rather than simple adoption. Mackie and Aspenlieder's (2024) rubric for evaluating AI tools offers fundamental criteria that can help reviewers organize evaluation around issues such as purpose, privacy, accessibility, equity, and tool limitations. In practice, frameworks like these give libraries a shared vocabulary for evaluation and make it easier to communicate why a tool raises concern beyond whether the tool appears useful in a demonstration.

Rubrics that are published, widely shared, and revised by the field can become important tools for institutions that are understaffed, underfunded, and underprepared, yet simultaneously asked to take on the workload of these evaluations. They also help define a shared system of beliefs about what is important and what concerns librarians and library organizations should hold vendors accountable for. More broadly, well-defined evaluation criteria can support comparison, identify areas for improvement, and make organizational claims about responsible AI more open to scrutiny (Alaga et al., 2024). However, rubrics are not tools of infinite value. There is always a temptation to create just one more rubric to clarify one more point, but this can lead to an overproliferation of rubrics. Continued reinvention can lead to diminishing returns as professional effort becomes stuck at rubric writing instead of moving toward the next steps. This sense of stagnation is due to the limits of what these tools can meaningfully capture about AI systems. Rubrics function on several key assumptions: that evaluators have accurate information, that product features are stable and well documented, and that test environments represent actual use (Alaga et al., 2024). In practice, librarians are often asked to evaluate vendor documentation that may be too vague to qualify as genuine technical specifications, may only partially explain how systems function once deployed, or may be tied to carefully constructed demo environments designed to present idealized results. Because of these factors, rubrics risk shifting effort toward scoring claims instead of digging deeper into the realities of how these tools function in live library environments.

We risk a compounded persistent lack of transparency from vendors themselves. Features can change after rollout without clear communication, sales representatives may be unable to answer technical questions, and privacy claims may read as aspirational rather than operational. When vendors are secretive about how tools function or deflect questions about their design, data, limitations, or vulnerabilities, libraries may be unable to distinguish legitimate technical constraints from a lack of meaningful disclosure (Nogueira et al., 2025). Promises can exceed technical reality and be replaced by vague language where specificity should be. A scoring framework cannot tell a library what happens to user data once it enters a vendor system, or what will change as these models evolve after implementation unless vendors surface that information. When the information needed to

evaluate a tool is unavailable or obscured, the rubric becomes a document of good intentions rather than a meaningful accountability mechanism.

Issues with scope present an equally significant challenge. Rubrics can only evaluate one item at a time, not the whole ecosystem, and libraries risk losing sight of the broader landscape. Libraries are not adopting tools in isolation; they are bringing tools into complex environments shaped by existing platforms, licensing structures, instruction workflows, access expectations, and user needs. Focusing on one product at a time and assigning scores to various attributes can obscure several important considerations: ongoing professional accountability to users even when systems are externally controlled, the invisible labor involved in ensuring products are maintained and updated, and the broader shift in the field as the accumulation of these products creates structural changes in library environments. Rubrics can answer for a library, "Is this tool acceptable?" but they cannot fully answer, "What is this pattern of tool adoption doing to our work and our responsibilities over time?"

AI Innovation or another AI Chatbot?

In our experience attending vendor demonstrations, AI-enabled library products are often presented as innovative without a clear explanation of how they differ from widely available platforms such as ChatGPT, Copilot, Gemini, or Claude. Vendors may describe new research assistants in broad terms while providing limited evidence of distinct functionality or improved outcomes. For academic libraries, this creates a specific evaluation problem: how can libraries determine whether a vendor AI research assistant meaningfully improves research, discovery, or learning when it appears to replicate general chatbot behavior?

In vendor demonstrations and rollout initiatives, libraries may be encouraged to evaluate AI-enabled features while those products and their governance expectations are still evolving. This positions library staff as early evaluators responsible for identifying limitations, troubleshooting workflows, and communicating concerns back to vendors. For example, Ex Libris introduced Primo Research Assistant as an optional AI-enabled feature available within Primo environments (Clarivate, 2024). Libraries using Primo must then determine whether and how the tool aligns with local instructional, privacy, accessibility, and service expectations. Even when these features are optional, their integration into widely used library systems can create institutional pressure to assess, explain, and potentially support tools whose long-term functionality and implications are still developing.

In an educational space already competing with other high-priced resources, vendors need to clearly explain the value coming from model fine-tuning, retrieval design, metadata structuring, content coverage, citation verification, or workflow integration. Marketing claims that AI will transform research are insufficient when vendors do not explain how AI is being used, why it adds value, or whether that value is proportionate to the tool's financial and less visible costs (Nogueira et al., 2025). Too often, the explanation stops at buzzwords, and AI becomes a label rather than a measurable improvement. Academic libraries should be cautious when vendor-embedded AI features are presented as innovative and ready to use without clear explanation of the testing environment, testing duration, sample size, privacy practices, data retention, source content, accuracy limits, accessibility, user support, and post-launch maintenance.

The result is a familiar pattern: heavy marketing language, limited technical clarity, and an expectation that librarians will fill the gaps for users. Libraries are on the receiving end of pressure from vendors operating under their own market pressures. A genuinely collaborative relationship, where both vendors and libraries function as partners, would be more likely to create products that suit library needs. If vendors want libraries to pilot, purchase, enable, or support these tools, then vendors should also be prepared to explain why their tools are meaningfully better than freely available chatbot

platforms and how their products align with the library's role in supporting research, teaching, and access.

The Risks of an AI Bubble

The rapid expansion of commercial and industry investment in AI following the release of ChatGPT prompted concerns that investment may be outpacing sustainable returns (Storm, 2025). Much like the dotcom bubble of 2000, observers of AI technologies do not generally believe the technology itself will fail or disappear, but rather that current enthusiasm has outpaced sustainable returns (Potrykus, 2024). Understanding how this market dynamic influences pricing, functionality, and adoption of decisions is becoming an essential part of AI evaluation. For libraries that have moved quickly to adopt AI-driven tools and services, the bubble dynamic introduces a distinct set of vulnerabilities. Institutions that have built workflows or allocated budget lines around specific commercial AI products may find themselves at risk in the case of pricing shifts or the quiet discontinuation of features. Over-reliance on tools whose long-term viability is tied to speculative investment rather than stable returns also mirrors the kind of infrastructural fragility that the dotcom era made visible too late for many organizations. There is also a reputational dimension: libraries that position themselves as early adopters of AI solutions may face difficult questions from their communities if those solutions prove unsustainable. This is not an argument against adoption, but rather for strengthening evaluative literacy that encourages greater industry transparency, promotes values-driven decision-making, and supports more sustainable practices across library and academic communities. Libraries can strengthen their evaluative literacy by looking beyond whether a tool appears useful and instead by examining the technical and governance choices that shape how emerging AI operates within their campus. Long-standing open-source communities provide a useful comparison because they have extensive experience debating how technologies should be defined, documented, maintained, and held accountable.

What OpenSource AI Reveals about Vendor AI

Open-source communities offer academic libraries a model of critical engagement as AI features are introduced into databases, discovery systems, and research platforms. At the Southern California Linux Expo (SCaLE), an annual Linux, open-source, and free software conference held in the Los Angeles area, AI is discussed within a community that has long examined how technical terms, standards, and responsibilities should be defined, documented, and sustained rather than simply assumed. Stefano Maffulli, former Executive Director of the Open Source Initiative (OSI) and a frequent speaker at SCaLE, advanced this discussion by examining how AI complicates established definitions of open-source software. In traditional open-source software, openness generally means that people could inspect, modify, and redistribute the source code. He states that an AI model is shaped not only by its code but also by its training data, parameter weights, hardware configuration, and deployment choices (Maffulli, 2023). For libraries, nuanced or inconsistent disclosure about these AI elements can create risk. The term "open" may imply transparency while still leaving libraries without enough clear information to understand how the system was developed, how it operates, or what limitations may affect its use. When vendor language remains open to interpretation, libraries may be left responsible for explaining and supporting claims they cannot confidently verify.

However, understandably, academic libraries cannot inspect proprietary AI systems in the same way that open-source communities can inspect source code, nor should librarians be expected to independently verify every technical component. They can, however, apply the same expectation that important claims should be clearly defined, supported by evidence, and documented. Libraries can ask vendors to explain terms such as "transparent," "trusted," "responsible," and "research-ready" in accessible language; provide documentation appropriate for institutional review; identify known limitations;

explain how product changes will be communicated; and clarify what technical support and accountability will remain with the vendor after implementation. Libraries can also ask vendors to clarify what technical support and accountability will remain with them after implementation. Practical questions might include: What information will remain available after the product launches? Where will stable documentation be maintained if product pages, announcements, or support links change? What evidence supports the vendor's claims? Who governs the tool's use, and what protections exist for users? When vendors cannot provide sufficient answers, libraries can document those gaps and consider them when deciding whether they are prepared to adopt and support the product.

Therefore, Maffulli's concern about the instability of "open source" as a term in AI has relevance for vendor-embedded AI in academic libraries. In traditional open-source communities, shared principles generally preceded the licenses created to protect them; the ethos came before the product. In the current AI environment, workflow and announcements from vendors often appears reversed. Products are being released, licensed, and commercialized while expectations for openness, accountability, and responsibility are still developing. Academic libraries experience this reversal when AI features are added to existing platforms, new licensing terms are introduced, and products are marketed before the field has reached a shared understanding of what responsible vendor-supported AI should require.

When capital and competition move faster than shared norms and governance structures, adoption can become market-driven rather than principle-driven. Institutions may feel pressured to move quickly because innovation is framed as urgency. For academic libraries, pressure is significant. Libraries are not simply consumers of technology; they are stewards of access, explanation, and long-term resource decisions. If foundational expectations for openness and accountability are still being established, academic libraries should approach adoption cautiously by participating actively in evaluation before making long-term commitments. For example, libraries can convene relevant stakeholders to assess a demonstration together, document concerns and desired changes, and ask the vendor to return with clearer information or a revised product before the library proceeds. An iterative process allows libraries to build AI literacy, define local principles, evaluate vendor claims, and assess potential impacts while there is still an opportunity to influence implementation.

The Most Valuable AI Investment May Not Be a Product for the Library

Libraries are in a unique position in the adoption and administration of AI technology. Professional experience in vendor evaluation, licensing negotiation, information literacy, and critical assessment of information systems collectively functions as an applied form of evaluative expertise that is deeply relevant to vendor-embedded AI tools. The continual repetition of concerns and conversations about AI can be tiring, and that fatigue can make it harder to build thoughtful evaluation practices. But moving slowly is not the same thing as refusing to move, and carefulness is not resistance to innovation. It is recognition that professional accountability requires clearer norms than the current market pace allows.

While the industry's appetite for endless AI expansion seems infinite, library budgets are not. The value proposition for any AI product must be genuinely convincing. If a discovery platform or database already under subscription introduces a chatbot feature, it is worth evaluating whether the library needs another chatbot before moving on to evaluating the tool's specific strengths and limitations. Taking a broader view that accounts for the entire technology ecosystem rather than evaluating each tool in isolation allows libraries to identify gaps and redundancies in a way that a single vendor demo cannot.

Libraries can also facilitate broader conversations by developing more awareness of open-source alternatives, highlighting

infrastructure implications, and pushing for greater transparency in pricing models, data privacy, and post-launch support. This does not mean every library must build its own AI tool or reject vendor products entirely. The more sustainable investment may be in the people, policies, and practices that allow libraries to evaluate AI with clarity. Staff training, shared guidelines, local expertise, decision documentation, and sustainable support structures give libraries a stronger position in campus and vendor conversations before money, staff time, and user trust are committed to another tool.

The same professional skills that lead library professionals to approach vendors with critical inquiry and thoughtful receptiveness also help libraries recognize where AI may offer meaningful value. Current evidence suggests that AI is often effective when applied to structured tasks that transform, extract, or organize existing information rather than generate new knowledge. For instance, Demers (2025) demonstrated the use of LLMs and semantic technologies to extract structured data from published literature for systematic reviews. Emerging agentic AI tools likewise have the potential to interact across platforms, reducing repetitive manual entry that increases the likelihood of human error and supporting more integrated workflows. Although these applications remain in an early stage of development, they are being explored for tasks such as scheduling automation and collection weeding (Oladokun, 2026). Libraries can use their purchasing power and evaluative expertise to encourage vendors to develop AI tools that address clearly identified research and library needs and demonstrate their value through specific, assessable, and accountable criteria (Alaga et al., 2024).

The most valuable investment a library can make in AI right now may not be a product at all. Staff training, dedicated expertise, and institutional capacity to experiment with and create local solutions will outlast many outsourced alternatives. Librarians and library workers who understand how these systems work as a more lasting resource are than vendor tools controlled by external financial pressures. An in-between stance means refusing passive roles while remaining open to responsible experimentation. It means defining what the library needs, what the library expects, and what the library is willing to support before vendors, and campus pressures define those responsibilities for us.

Conclusion

Academic libraries contribute knowledge of research practices, source evaluation, information literacy, scholarly resources, and user needs after a tool becomes technically available. Our responsibility is not only to complete another rubric or watch another product demonstration or to be part of the AI conversation after the vendor has already shaped the feature, the timeline, or the terms of use. Our responsibility is to define what AI can meaningfully do for our users, whose research needs, access barriers, information practices, and support questions libraries evaluate every day. Therefore, Academic libraries should resist being positioned as passive recipients when it comes to opting into AI tools or features. Before endorsing, licensing, activating, or supporting AI-enabled features, libraries should ask what information our communities need, what problems the tool actually solves, what evidence supports vendor claims, and what conditions must be met for the tool to align with library values, user needs, and long-term resource responsibilities then demo until there is enough measurable return of investment.

AI directly affects core academic values that center originality, authorship, evidence, and the human voice. Students and faculty are already interacting with generative systems that raise new questions about intellectual labor, plagiarism, and the boundaries of research. In response, academic libraries will likely continue to participate in AI-enabled services that support campuswide strategic initiatives tied to student success, workforce development, and creative development. Participation, however, should not mean passive acceptance. Even when campus priorities or vendor roadmaps move ahead first,

libraries still need to respond with critical evaluation, measurable outcomes, and clear expectations for what these tools should do for users.

The most sustainable path is not simply to buy another AI product. Libraries need the capacity to decide how AI should be used, supported, explained, and revisited. That means defining local values, setting boundaries of use, clarifying privacy expectations, disclosing where AI is present, and naming who is accountable for updating guidance as tools change. It also means investing in staff training and shared confidence so that AI responsibility does not fall only on the few people who already understand the systems, licenses, platforms, or vendor relationships.

When capital and competition move faster than shared norms, adoption becomes market-driven rather than principle-driven. For academic libraries, this pressure is significant because libraries are not simply consumers of technology. Libraries are stewards of access, privacy, explanation, and long-term resource decisions. If vendors are already mapping out AI features for the platforms that libraries rely on, libraries should be more direct as teams, cohorts, or as part of professional communities about how the user wants to use AI responsibly.

Remaining critically in-between offers a more sustainable path. In a system like the CSU, a proactive approach could include shared vendor-evaluation guidelines, campus representation, required training, and local policies that clarify how AI tools should be assessed and supported. At the campus level, libraries should advocate for dedicated expertise, actionable training, and a living AI roadmap before investing heavily in AI products. Libraries need to push back when AI adoption focuses only on purchasing tools rather than building the knowledge, staffing, and evaluation practices needed to support those tools responsibly.

Optimism does not have to mean blind adoption. Libraries can acknowledge that many tools are still unpolished while remaining curious about how the technology may responsibly develop. A hopeful and critical approach requires asking whether a tool genuinely advances the library's mission, whether its infrastructure aligns with library values, and whether the library is prepared to support the tool over time. Most importantly, an in-between stance affirms that libraries are not passive recipients of AI integration. The most sustainable approach is for academic libraries to participate actively as the field unfolds, integrating AI in ways that reflect library priorities, institutional values, and long-term accountability.

References

About SCALE | SCALE. (n.d.). https://www.socallinuxexpo.org/scale/23x/about-scale

Alaga, J., Schuett, J., & Anderljung, M. (2024). A Grading Rubric for AI Safety Frameworks (Version 1). arXiv. https://doi.org/10.48550/arxiv.2409.08751

Caico, M., Harris, L., & O'Shea, S. (2024). Is This AI Tool Right for Me? Important Questions from the Framework. College & Research Libraries News, 85(8), 341. https://doi.org/10.5860/crln.85.8.341

California State University. (2025, February 04). CSU announces landmark initiative to become nation's first and largest AI-empowered university system. Retrieved February 13, 2026, from https://www.calstate.edu/csu-system/news/Pages/CSU-AI-Powered-Initiative.aspx

Clarivate. (2024, September 18). Clarivate Launches Generative AI-Powered Primo Research Assistant. Retrieved May 14, 2026, from https://www.prnewswire.com/news-releases/clarivate-launches-generative-ai-powered-primo-research-assistant-302250323.html

Demers, C. (2025). Applying LLMs and Semantic Technologies for Data Extraction in

Literature Reviews. Proceedings of the Annual Conference of CAIS / Actes Du Congrès Annuel de l'ACSI. https://doi.org/10.29173/cais2020

Industry Source. (2025, December 11). Universities surge in AI majors. Technical Education Post. https://www.techedmagazine.com/universities-surge-in-ai-majors/

Mackie, K., & Aspenlieder, E. (2024). Rubric for evaluating AI tools: Fundamental criteria. Paul R. MacPherson Institute for Leadership, Innovation and Excellence in Teaching, McMaster University.

Maffulli, S. (2023, March 11). Defining an open source AI [Conference presentation]. SCALE 20X, Ballroom F, Pasadena, CA, United States. https://www.socallinuxexpo.org/scale/20x/presentations/defining-open-source-ai/

McDonald, N., Johri, A., Ali, A., & Collier, A. H. (2025). Generative artificial intelligence in higher education: Evidence from an analysis of institutional policies and guidelines. Computers in Human Behavior: Artificial Humans, 3, Article 100121. https://doi.org/10.1016/j.chbah.2025.100121

Nogueira, L., Moltubakk, S., Fagervik, A., & Langfeldt, I. (2025). Cutting through the noise: Assessing tools that employ artificial intelligence. IFLA Journal, 51(3), 735–749. https://doi.org/10.1177/03400352241304121

Oladokun BD (2026), "Agentic AI and the future of library services". Library Hi Tech News, Vol. 43 No. 3 pp. 48–51, https://doi.org/10.1108/LHTN-11-2025-0199

Potrykus, M. (2024). Dot-com and AI bubbles: Can data from the past be helpful to match the price bubble euphoria phase using dynamic time warping? Finance Research Letters, 67, 105799. https://doi.org/10.1016/j.frl.2024.105799

Storm, S. (2025). The AI bubble and the U.S. economy: How long do "hallucinations" last? (Working Paper No. 240). Institute for New Economic Thinking. https://doi.org/10.36687/inetwp240

Validated Insights Inc. (2025) The demand for AI programs: Navigating the AI education landscape. https://viaieducation.carrd.co/