GCI Response to HHS Request for Information: Accelerating the Adoption and Use of Artificial Intelligence as Part of Clinical Care (RIN 0955-AA13)

Submitted electronically to Regulations.gov
February 23, 2026

Executive Summary

The General Catalyst Institute (GCI) appreciates the opportunity to respond to the Department of Health and Human Services' Request for Information regarding the acceleration of artificial intelligence adoption in clinical care. As an organization dedicated to promoting market-driven solutions powered by applied AI to transform healthcare delivery, we believe our insights and the experiences of our portfolio companies can provide valuable perspectives on regulatory modernization, reimbursement reform, and research priorities that will position the United States as the global leader in healthcare AI innovation.

We are encouraged by HHS' "OneHHS" AI Strategy and the unprecedented alignment across federal agencies to harness the transformative potential of AI while safeguarding patient privacy and civil liberties. The United States stands at a critical inflection point where strategic policy decisions made today will define our global leadership and reinforce the resilience of our health systems for decades to come. By addressing the critical levers of regulation, reimbursement, and research & development, HHS can reduce the uncertainty that currently impedes innovation and align federal incentives to enhance productivity, reduce burden, lower healthcare costs, and improve health outcomes.

Our message is simple: the tools exist. The talent exists. The policy pathways must now align to ensure these assets can scale equitably and sustainably across every level of the healthcare system.

Company Information

The General Catalyst Institute was launched to promote and strengthen national resilience by backing transformative technologies and shaping public policies that improve society. Our mission is to cultivate a healthy ecosystem for entrepreneurship and serve as a partner to policymakers on technology adoption, with a particular emphasis on applied AI in healthcare.

Central to this commitment to long-term transformation is our health assurance mission to create a more affordable, accessible, proactive, and equitable system of care by leveraging innovative technology, new business and payment models, and human capital. We execute this vision through a unique three-pronged approach: investing in innovative portfolio companies,

partnering with 27+ health systems and payers, and operating the Health Assurance Transformation Company (HATCo) as a blueprint for systemic change.

Our Catalyzing Care: A Framework for a Healthier America, introduced in the "U.S. Healthcare That Works" report, outlines five pillars for healthcare transformation:

1. Foster Healthier Outcomes for All
2. Refine Needed Innovations Without Red Tape
3. Advance Patient-First Care, Data and Agency
4. Maximize Fiscal Responsibility for U.S. Healthcare
5. Enhance U.S. Medical Talent for Today and Tomorrow

Each of these pillars has associated short-term and long-term policy recommendations that directly inform our response to this RFI. If implemented, the impact for real people would be remarkable: doctors could save thousands more lives each year, patients in rural and underserved communities could connect with specialists regardless of location through AI-enhanced care networks, and hundreds of thousands of unnecessary hospitalizations could be prevented through proactive, AI-enabled interventions.

Bottom Line Up Front

The U.S. healthcare system stands at a historic inflection point. After decades of fragmentation, rising costs, and uneven outcomes, we now have the technological capabilities to fundamentally transform healthcare delivery through market-driven solutions powered by applied AI. The General Catalyst Institute believes HHS can accelerate this transformation by focusing on five critical areas where our ecosystem has developed valuable insights.

First, we have concrete examples of what's already working. Our portfolio includes companies demonstrating that AI can achieve clinical outcomes equivalent or superior to traditional care while reducing costs by 50-60%, reclaiming 15-20 hours per week of clinician time consumed by administrative burden, and enabling earlier detection of time-sensitive conditions that prevent hundreds of thousands of dollars per patient in downstream costs. These proven applications (spanning radiology analysis, clinical documentation, chronic disease management, remote patient monitoring, and generative AI for healthcare communications) are ready for widespread deployment but face regulatory uncertainty, reimbursement misalignment, and interoperability barriers that HHS can address.

Second, we can identify the specific blocking factors that impede scale. Through our portfolio companies' experiences deploying AI across diverse healthcare settings, we have documented where current regulations create ambiguity (such as unclear FDA oversight boundaries for clinical decision support tools versus medical devices), where payment models misalign with AI's value proposition (fee-for-service structures that don't reward prevention or efficiency gains), and where data fragmentation prevents AI development (siloed EHR systems requiring custom integration work that can take 12-18 months per health system). These barriers are solvable through targeted policy reforms we detail in our responses.

Third, we understand the incentive structures needed to drive adoption. Expanding value-based payment models creates natural demand for AI tools that improve outcomes and efficiency; when payers take financial risk for long-term health costs, they have inherent incentives to promote high-value AI interventions. Mandating interoperability standards and creating data commons would further reduce AI development costs by 40-60%, democratizing innovation beyond well-capitalized companies and enabling startups to compete on algorithmic merit rather than integration resources.

To capture these opportunities, HHS should foster competitive ecosystems of private sector innovation rather than mandating top-down solutions. This includes supporting industry-driven accreditation bodies that can move with the speed of innovation, creating regulatory safe harbors for proven AI applications, and investing in shared infrastructure (such as federated learning networks, benchmark datasets, and standardized testing environments) that reduces barriers to entry while maintaining rigorous oversight.

AI offers a fundamental shift from clinical scarcity (too few providers, too little time, too little access) to clinical abundance, where technology extends the reach of every clinician and makes high-quality care available to far more patients. By aligning regulation, reimbursement, and R&D investments, we can reduce the uncertainty that impedes innovation, foster public trust in modern technology, and ensure AI is deployed in ways that enhance productivity, lower costs, and improve outcomes for patients, caregivers, and communities across America. **T**he General Catalyst Institute and our ecosystem stand ready to partner with HHS, providing coordination support for pilots and testbeds and ongoing feedback to ensure HHS programs meet evolving technology and patient needs.

What follows is a table of our Detailed Responses to the HHS RFI. In addition, at the end of our response is a table of acronyms employed.

Detailed Responses

RFI Question	GCI Recommendation	Supporting Evidence	Expected Impact
Question 1: **What are the biggest barriers to private sector innovation in AI for health care and its adoption and use in clinical care?**	HHS should address three interconnected barriers: **(1) Regulatory uncertainty around which AI applications require FDA oversight versus clinical decision support exemptions.** **(2) Misaligned reimbursement incentives where fee-for-service models don't reward AI tools that prevent costly interventions or improve efficiency.** **(3) Fragmented data infrastructure requiring custom integration for each deployment.** Specifically, HHS should clarify FDA regulatory pathways under 21 CFR Part 860 for different AI risk categories, expand value-based payment models under 42 CFR Parts 414 and 425 that reward outcomes rather than volume, and mandate comprehensive FHIR-based interoperability under 45 CFR Part 170 to eliminate data silos.	AI-powered radiology solutions for time-sensitive conditions like strokes and pulmonary embolisms have cut time-to-treatment by more than 60% in clinical studies, directly averting permanent disability and the hundreds of thousands of dollars in long-term care costs that delayed diagnosis produces. Yet reimbursement structures don't reward this efficiency. Current payment models create perverse incentives: faster triage reduces billable imaging, while preventing complications eliminates readmissions. Adoption has been slowed by separate contract negotiations and integration with dozens of PACS systems. Conversational AI platforms face regulatory uncertainty about classification and evidence standards. Despite investment in safety testing and bias mitigation, companies lack guidance on FDA premarket review requirements, post-market surveillance structure, or liability frameworks. This uncertainty discourages investor confidence and creates a chilling effect on patient-facing AI innovation. Healthcare operating system companies (those building the infrastructure layer that enables AI	Regulatory clarity would enable companies to plan development roadmaps with confidence, attracting capital for long-term R&D. Analysis estimates regulatory streamlining for low-risk AI could accelerate time-to-market by 18-24 months, bringing life-saving technologies to patients years earlier. Reimbursement reform aligning payment with outcomes would create market pull for AI tools that genuinely improve care efficiency, driving investment toward high-value applications. This shift is essential to capture the deflationary benefits of AI. Current fee-for-service models are inflationary by nature; conversely, value-based models allow the system to absorb AI's productivity gains, passing the savings generated by automation back into the broader healthcare ecosystem. Mandated interoperability would reduce deployment costs by 40-60%, making AI economically viable for smaller health systems and safety-net hospitals, democratizing access across all patient populations. Together, these changes align with our Refine Needed Innovations Without Red Tape pillar. This illustrates the deflationary power of proactive AI: it shifts spending from expensive, reactive acute care toward early, low-cost algorithmic intervention, effectively 'compressing' the total cost of a patient's care journey.

RFI Question	GCI Recommendation	Supporting Evidence	Expected Impact
		deployment at the point of care) encounter significant data fragmentation when connecting to disparate healthcare systems. Basic data exchange requires navigating proprietary interfaces, custom HL7 implementations, and restrictive governance policies. Without mandatory interoperability standards, AI companies must build custom integrations for each provider (which can mean building dozens of integrations) with each implementation taking 12–18 months to access the real-world clinical data needed to train and validate AI tools.	
Question 2: **What regulatory, payment policy, or programmatic design changes should HHS prioritize to incentivize the effective use of AI in clinical care and why?** **What HHS regulations,**	HHS should prioritize three core policy reforms: **(1) Tiered Regulatory Framework:** Modernize 21 CFR Part 860 creating risk-based AI clinical decision support categories. Exempt low-risk AI from premarket review while maintaining post-market surveillance. Medium-risk applications should follow streamlined 510(k) pathway. **(2) Outcome-Based Reimbursement:** Establish payment mechanisms under 42 CFR Parts 414 and 425 recognizing AI-enabled care and	(1) Tiered Regulatory Framework: The absence of risk-stratified AI oversight creates two compounding problems. First, lower-risk AI clinical decision support tools face the same premarket review burden as higher-risk devices, creating cost and timeline barriers that deter development and deployment, particularly for smaller companies. Second, existing postmarket surveillance requirements were not designed for AI systems that update continuously, leaving regulators without adequate tools to monitor algorithmic updates in deployed products. A tiered framework would calibrate premarket requirements to actual risk while establishing AI-specific postmarket surveillance obligations that protect	Tiered regulatory framework would reduce time-to-market for low-risk AI by 2-3 years while maintaining safety oversight, potentially accelerating patient access to hundreds of beneficial applications currently stalled in regulatory uncertainty. Outcome-based payment models would redirect an estimated $50-100 billion in current healthcare spending toward high-value AI interventions, creating sustainable revenue streams and enabling health systems to invest in AI infrastructure without immediate margin pressure. Interoperability mandates and data commons (data repositories) would reduce AI development costs by 40-60% by eliminating redundant integration work and enabling training on representative datasets, democratizing AI development beyond well-capitalized companies.

RFI Question	GCI Recommendation	Supporting Evidence	Expected Impact
policies, or programs could be revisited to augment your ability to develop or use AI in clinical care? **Please provide specific changes and applicable Code of Federal Regulations citations.**	the distinction between AI-only tools and AI-supported services. Expand value-based models rewarding outcomes over volume. Allow ACOs to share AI savings with vendors and create Stark Law safe harbors for AI tools improving care quality. **(3) Accelerated Interoperability:** Mandate FHIR-based APIs for certified EHRs and prohibit information blocking under 45 CFR Part 170. Create NIH data commons for qualified AI developers. Establish clear HIPAA guidance on de-identified datasets for AI development.	patients without impeding innovation. (2) Outcome-Based Reimbursement: Current payment structures actively penalize AI-enabled innovation. Psychotherapy reimbursement codes do not recognize AI-assisted therapy planning, making adoption economically unviable even where AI-enhanced treatment matching improves outcomes. Stark Law exposure under 42 CFR §411.357 similarly deters health systems from deploying AI-powered care coordination tools to referring physicians, even when those tools reduce unnecessary utilization. Safe harbors and updated payment mechanisms under 42 CFR Parts 414 and 425 would directly remove these barriers. (3) Accelerated Interoperability: AI tools are only as effective as the data they can access. Fragmented, proprietary data environments (compounded by the absence of clear HIPAA guidance on de-identified datasets) slow development and increase legal risk, particularly for smaller companies. Mandated FHIR-based APIs, enforceable information blocking prohibitions under 45 CFR Part 170, and a NIH data commons would accelerate the evidence generation necessary to support coverage and	Together, these reforms align with our Maximize Fiscal Responsibility and Advance Patient-First Care, Data and Accessibility pillars, positioning the United States as global leader in healthcare AI innovation while ensuring benefits accrue broadly across all demographics and care settings.

RFI Question	GCI Recommendation	Supporting Evidence	Expected Impact
		payment reform.	
Question 3: **For non-medical devices, we understand that use of AI in clinical care may raise novel legal and implementation issues that challenge existing governance and accountability structures (e.g., relating to liability, indemnification, privacy, and security).** **What novel legal and implementation issues exist and what role, if any, should**	HHS should convene multi-stakeholder frameworks addressing three critical legal issues: **(1) Distributed Liability:** Establish a working group including AI developers, healthcare providers, malpractice insurers, patient advocates, and legal experts to develop model liability allocation frameworks. Frameworks should clarify responsibility based on AI autonomy level, clinical context, and decision nature. **(2) Algorithmic Transparency:** Work with NIST and ONC to establish tiered explainability standards appropriate to clinical context and risk level. Life-or-death decisions might require deep algorithmic transparency with trade secret protections, while lower-risk administrative AI could satisfy requirements through input-output explanations. **(3) Consent Frameworks:** Clarify HIPAA requirements under 45 CFR §164.508 for AI use disclosure and patient consent	Generative AI platforms for patient-facing healthcare applications face fundamental uncertainty about liability allocation. If an AI agent provides guidance resulting in an adverse outcome, is the ultimate liability with the developer, healthcare organization, overseeing clinician, or patient? Traditional liability and agency law don't map cleanly onto semi-autonomous AI systems, forcing highly customized indemnification provisions that increase contracting complexity. Radiology AI platforms illustrate the transparency challenge. Radiologists want to understand why AI flagged certain cases for clinical decision-making and medicolegal protection. However, full transparency into neural network architecture could enable competitors to reverse-engineer R&D investment. Companies have developed explainability features showing which image features AI weighted most heavily, but legal uncertainty remains about whether this satisfies disclosure obligations in litigation. Healthcare infrastructure platforms encounter consent challenges regularly. When systems use AI to suggest	Clear liability frameworks would reduce contracting friction, enable more transparent adverse event reporting, and accelerate adoption of beneficial AI tools by reducing legal risk uncertainty. Tiered explainability standards would enable AI developers to build consistent transparency mechanisms rather than negotiating bespoke solutions with each health system, creating a more efficient market for clinical AI and reduced legal costs for all stakeholders. Clear HIPAA requirements for AI disclosure and consent would enable consistent patient communication informed by actual clinical risk rather than legal uncertainty, building greater patient trust. These frameworks align with our Advance Patient-First Care, Data and Accessibility pillar by ensuring patients understand and can make informed decisions about AI involvement in their care while enabling innovation to proceed with appropriate safeguards.

RFI Question	GCI Recommendation	Supporting Evidence	Expected Impact
HHS play to help address them?	through rulemaking establishing safe harbors. Address patient consent requirements for AI analysis, opt-out rights, and continuous learning systems.	scheduling, flag drug interactions, or predict no-show risk, these analyses influence care but may not be visible to patients. Current regulations don't clearly establish whether behind-the-scenes AI use requires explicit patient notification or consent. This is particularly acute for continuous learning systems where AI models evolve based on aggregated clinical data without patient awareness or ability to re-consent.	
Question 4: **For non-medical devices, what are the most promising AI evaluation methods (pre- and post-deployment), metrics, robustness testing, and other workflow and human-centered evaluation**	HHS can support evaluation through: **(1) Standardized Frameworks:** HHS should establish contracts for developing standardized evaluation frameworks addressing technical evaluation (algorithmic performance) and implementation evaluation (workflow integration). **(2) Pragmatic Trials:** NIH should fund cooperative agreements partnering AI companies with academic medical centers to conduct pragmatic clinical trials embedded within real workflows, measuring algorithmic accuracy and implementation outcomes. **(3) Continuous Monitoring:**	AI-enabled physical therapy platforms have conducted pragmatic trials randomizing patients to AI-enabled care versus standard treatment within real-world insurance populations. Trials revealed AI effectiveness varied significantly based on patient engagement levels and technology literacy. Based on findings, platforms modified onboarding and developed predictive models to identify patients needing additional human support. This iterative evaluation enabled continuous improvement but was funded entirely through private capital because no clear federal funding mechanisms exist for post-deployment evaluation of non-device AI tools. Radiology AI platforms have implemented sophisticated	Standardized evaluation frameworks would reduce duplicative effort and establish common performance benchmarks, enabling more efficient assessment of AI tools. Pragmatic trials through cooperative agreements would address market failure where individual companies lack incentives to share evaluation data, even though collective knowledge advancement would accelerate the field. CMS demonstrations would provide evidence for coverage and reimbursement decisions while generating methodological insights. Interoperable monitoring infrastructure would democratize post-deployment surveillance capabilities, enabling smaller companies to maintain rigorous safety oversight. Prize competitions would catalyze innovation in evaluation science itself.

RFI Question	GCI Recommendation	Supporting Evidence	Expected Impact
methods for clinical care? **Should HHS further support these processes?** **If so, which mechanisms would be most impactful (e.g., contracts, grants, cooperative agreements, and/or prize competitions)?**	ASTP/ONC and AHRQ should fund grants for developing interoperable AI monitoring infrastructure embedded within EHR systems, enabling automated performance tracking of technical metrics and clinical outcomes. **(4) Demonstration Projects:** The CMS Innovation Center should establish demonstrations testing AI tools in Medicare/Medicaid populations while rigorously evaluating clinical and economic outcomes. **(5) Prize Competitions:** HHS should consider prizes for novel evaluation methods addressing current gaps.	post-deployment monitoring tracking AI performance across diverse hospital settings. They discovered AI accuracy for detecting intracranial hemorrhage varied significantly across CT scanner manufacturers and imaging protocols, a data drift issue only apparent through continuous real-world monitoring. Platforms now flag performance variations and work with radiology departments to adjust scanning protocols, but this requires substantial engineering resources smaller AI companies may not possess. Generative AI healthcare platforms are pioneering human-centered evaluation methods including structured testing of how patients and clinicians respond to AI-generated communications. They've found seemingly small variations in AI communication style dramatically affect patient trust and adherence. However, developing rigorous evaluation methods for conversational AI requires interdisciplinary expertise that most companies cannot sustain.	The expected outcome is a robust evidence base guiding AI adoption decisions, identifying best practices, and building public and clinician trust in AI-enabled care. This infrastructure aligns with our Foster Healthier Outcomes for All pillar by ensuring AI tools are rigorously evaluated for safety, effectiveness, and equity.
Question 5: **How can HHS best support private sector activities**	HHS can incentivize voluntary industry-driven standards while providing federal recognition and safe harbor provisions: **(1) Deemed Status Model:** Designate qualified accreditation organizations (similar to hospital	Currently each health system conducts its own ad hoc evaluation of AI tools, requiring companies to undergo duplicative vendor assessment processes examining similar issues (data security, algorithmic bias, workflow integration) with different methodologies and documentation	By supporting private sector certification and accreditation, HHS would create a more efficient and trustworthy AI healthcare marketplace. Healthcare organizations could confidently adopt certified AI tools knowing they meet rigorous standards, reducing the need for extensive internal evaluation that many organizations lack capacity to perform. This would particularly benefit smaller community

RFI Question	GCI Recommendation	Supporting Evidence	Expected Impact
(e.g., accreditation, certification, industry-driven testing, and credentialing) to promote innovative and effective AI use in clinical care?	accreditation under 42 CFR §488.5) that can certify AI systems meeting established safety, effectiveness, and equity standards. Healthcare organizations using HHS-recognized certified AI tools would receive regulatory advantages. **(2) Consensus Standards:** Convene stakeholders to establish consensus around core principles for AI certification. Work with NIST, National Quality Forum, and AI-specific organizations to develop outcomes-focused, technology-neutral frameworks. **(3) Shared Testing Infrastructure:** Support development of industry-driven testing infrastructure to include shared testbeds, benchmark datasets, and testing environments that preserve patient privacy. **(4) Clinical AI Credentialing:** Support credentialing programs training healthcare professionals to effectively evaluate, implement, and oversee AI systems.	requirements. If HHS recognized industry-driven certification that health systems could rely upon, companies could complete a single rigorous accreditation process rather than hundreds of individual vendor reviews, accelerating market entry particularly for smaller companies lacking dedicated regulatory affairs teams. Healthcare infrastructure companies face challenges validating AI-enabled workflows because each health system has different EHR configurations, data quality levels, and clinical processes. Creating synthetic yet realistic testing environments where AI tools could be validated against standardized clinical scenarios would significantly reduce deployment risk. HHS-supported shared testing infrastructure would level the playing field and accelerate innovation. AI tools for treatment matching and outcome prediction require clinical oversight, but most mental health professionals lack training in AI literacy (understanding what AI can do, when to trust AI recommendations, how to identify potential algorithmic bias). Industry-wide credentialing standards would create portable skills clinicians could apply across different AI tools and care settings, building clinical workforce capacity needed for	hospitals and rural health systems, democratizing access to AI-powered care improvements beyond well-resourced academic medical centers. For AI developers, certification would provide valuable market differentiation and reduce sales cycles, enabling faster scaling of proven technologies. Designation of HHS-recognized accreditation bodies would create healthy competition among certification organizations to develop most rigorous and clinically meaningful standards. Shared testing infrastructure would reduce AI development costs by 20-30%, enabling more companies to validate tools across diverse populations before deployment, improving overall AI quality and safety while accelerating innovation. Credentialing for clinical AI specialists would create career pathways attracting talent to this emerging field and ensure AI adoption is clinician-led.

RFI Question	GCI Recommendation	Supporting Evidence	Expected Impact
		widespread AI adoption while ensuring appropriate professional oversight.	
Question 6: **Where have AI tools deployed in clinical care met or exceeded performance and cost expectations and where have they fallen short?** **What kinds of novel AI tools would have the greatest potential to improve health care outcomes, give new insights on quality, and help reduce costs?**	AI tools have exceeded expectations in well-defined clinical domains, such as diagnostic imaging and clinical documentation. Many AI tools and solutions have been commercially available for some time and the use case has proven itself across market segments. The greatest current and future potential lies in three categories: **(1) Generative AI Agents:** Automate administrative burden, enabling clinicians to focus on patient care. AI agents handling routine communications, authorizations, and documentation could reclaim 15-20 hours per week of clinician time, addressing burnout. **(2) Multimodal AI:** Integrate diverse data streams for holistic patient understanding. Systems synthesizing information from EHRs, imaging, genomics, and social determinants could identify patterns humans cannot detect, enabling earlier intervention. **(3) Predictive AI:** Enable proactive rather than reactive	AI-enabled physical therapy, radiology AI, and clinical documentation illustrate where AI has delivered. Randomized controlled trials demonstrate AI-guided physical therapy achieves equivalent outcomes to in-person care at approximately 60% lower cost, with one platform now serving 2,500+ enterprise clients and 400,000+ members. Radiology algorithms achieve sensitivity above 90%, though hospitals that redesigned workflows around AI capabilities saw dramatically better time-to-treatment outcomes than those that did not, underscoring that algorithmic performance alone is insufficient. Generative AI agents have demonstrated meaningful capacity to handle routine clinical communications, though early deployments revealed that patient acceptance depends heavily on empathy and contextual understanding. Companies that invested in naturalistic dialogue models have seen stronger adoption, pointing toward significant potential to reclaim clinician time if implementation prioritizes human-centered design. Multimodal and predictive AI represent the next frontier: systems synthesizing EHR, imaging, genomics, and social	Near-term, HHS should prioritize expanding access to AI tools with robust evidence of their efficacy (such as radiology AI, digital therapeutics, remote monitoring) to underserved populations. Could prevent 50,000-100,000 adverse events annually and reduce costs by $5-10 billion. Democratizing AI-enhanced radiology could reduce stroke care disparities. Medium-term, HHS should catalyze generative AI agents automating administrative burden, potentially reclaiming 1-2 billion hours annually of clinician time. This would address the workforce crisis while reducing administrative costs, estimated at $250-400 billion annually. Long-term, HHS should invest in multimodal AI and predictive models enabling personalized, proactive medicine. Could identify at-risk patients years before symptoms emerge, shifting spending from acute care toward prevention and early intervention. The ultimate goal is a deflationary shift in the national health spend: moving from a system that 'manages' rising costs to one where AI-driven productivity and preventive care consistently drive the real cost of quality outcomes downward.

RFI Question	GCI Recommendation	Supporting Evidence	Expected Impact
	care. Models identifying patients likely to deteriorate could prevent 30-40% of hospitalizations through timely intervention. AI has fallen short when implementation challenges (such as alert fatigue, poor workflow integration) undermine algorithmic performance.On cost, point solutions frequently optimize individual workflows without producing system-level savings which is a reflection of fragmented implementation rather than AI's inherent limitations.	determinants data are beginning to identify deterioration signals and intervention opportunities at a scale and speed no manual care management program can replicate. Underperformance has been driven by implementation failures rather than algorithmic limitations. Alert fatigue, poor workflow integration, and deployment outside training environments consistently undermine tools that perform well on benchmarks. Cost expectations have similarly gone unmet when point solutions optimize individual workflows without producing system-level savings.	
Question 7: **Which role(s), decision maker(s), or governing bodies within health care organizations have the most influence on the adoption of AI for clinical care?**	AI adoption requires alignment across multiple stakeholders with often misaligned incentives: **(1) Key Decision Makers:** Chief Medical Information Officers evaluate clinical utility; Chief Information Officers assess technical feasibility and security; Chief Financial Officers determine economic viability. CMIOs prioritize clinical effectiveness, CIOs focus on interoperability and cybersecurity, CFOs evaluate ROI using traditional frameworks that struggle to capture AI's value proposition.	Through partnerships with 27+ health systems, we have observed that successful AI adoption almost always requires a champion who advocates for AI tools with both clinical leadership and medical staff. Without clinical champions, AI initiatives stall regardless of technical merit. Several health systems initially rejected radiology AI proposals because radiologists expressed concerns about losing autonomy. Successful adoption occurred when radiology department leaders reframed AI as a tool to enhance radiologist capabilities and when early adopters demonstrated benefits to skeptical colleagues.	HHS can address adoption through: Standardized Evaluation Frameworks: HHS could create playbooks guiding health systems through AI assessment, including template evaluation criteria and implementation checklists. Would particularly benefit smaller community hospitals lacking dedicated innovation teams. Expand Value-Based Payment: CMS should expand value-based payment models aligning financial incentives with AI adoption. Additionally, CMS could create specific quality measures incentivizing AI adoption for evidence-based applications. Change Management Support: HHS should support change management and workforce training

RFI Question	GCI Recommendation	Supporting Evidence	Expected Impact
What are the primary administrative hurdles to the adoption of AI in clinical care?	**(2) Governance Gaps:** Many health systems lack clear processes for evaluating AI tools, leading to ad hoc decision-making. Hospital committees often lack AI expertise, leading to either excessive caution or insufficient scrutiny. Medical staff credentialing wasn't designed for AI-enabled care, creating ambiguity about approval requirements. **(3) Primary Administrative Hurdles:** Procurement complexity (lengthy RFP processes, legal negotiations, approval layers); inadequate IT infrastructure (legacy EHR systems lacking modern APIs); organizational culture and clinician resistance (skepticism, liability concerns, workflow resistance).	CFO approval represents a critical bottleneck, particularly for AI tools improving quality without generating new revenue. In fee-for-service payment models, CFOs struggle to justify investments in AI that might reduce billable services. This misalignment is particularly acute for AI applications in population health and preventive care. In contrast, health systems with substantial value-based payment contracts demonstrate much greater willingness to invest in AI tools improving outcomes, even without immediate margin impact. Healthcare infrastructure companies have observed health systems often default to AI tools embedded within existing EHR systems even when third-party solutions demonstrate superior performance, because integrated solutions avoid separate procurement processes and reduce IT integration complexity. This "bundling advantage" for incumbent EHR vendors creates significant barriers to entry for specialized AI companies.	initiatives building organizational capacity for AI adoption, including funding collaboratives where health systems share AI implementation best practices. Expected impact is a reduction in a 12–18 month adoption timeline to 3-6 months for health systems with streamlined evaluation processes. This acceleration would enable patients to benefit from AI innovations years earlier. Aligns with our Enhance U.S. Medical Talent pillar by building workforce capacity to effectively leverage AI tools.
Question 8: **Where would enhanced interoperabili**	Enhanced interoperability would accelerate development in: **(1) Longitudinal Clinical Data Integration:** Mandate	Healthcare infrastructure companies reveal complexity of current interoperability challenges. Even when health systems claim to have FHIR APIs implemented, actual data	Enhanced interoperability would reduce AI development costs by estimated 40-60% by eliminating redundant data integration work, enabling AI companies to focus on algorithmic innovation. Would particularly benefit smaller

RFI Question	GCI Recommendation	Supporting Evidence	Expected Impact
ty widen market opportunities, fuel research, and accelerate the development of AI for clinical care? **Please consider specific data types, data standards, and benchmarking tools.**	FHIR-based data exchange across all care settings. Currently AI developers struggle to access complete patient journeys because data remains fragmented across hospitals, primary care, specialists, pharmacies, payers, and patient-generated sources. **(2) Multimodal Data Types:** ASTP/ONC should mandate comprehensive FHIR implementation including standardized approaches for exchanging imaging data, clinical notes, genomic data, and patient-generated health data from apps and wearables. Standards needed for exchanging AI model outputs so downstream AI systems and clinicians can access predictions. **(3) Federated Learning Infrastructure:** HHS should invest in technical infrastructure and governance frameworks enabling federated learning at scale. Federated learning trains AI models by sending algorithms to data at each institution, aggregating only model updates rather than raw patient information, enabling training on millions of patients while respecting HIPAA requirements.	availability varies dramatically. Inconsistent FHIR implementation creates "lowest common denominator" problems where AI developers can only rely on universally available data elements, which often excludes most clinically valuable information. Deploying AI-enabled workflows requires 3-5x more engineering effort for data integration than for algorithm development itself. Radiology AI algorithms for detecting stroke could be significantly improved by incorporating non-imaging data like patient vital signs and medication history, but accessing this information requires custom integration with each hospital's EHR system. If FHIR-based interoperability enabled seamless access to relevant clinical context alongside imaging data, AI models could provide more accurate predictions while requiring far less custom integration work. AI-enabled physical therapy platforms capture detailed information about patient exercise performance and engagement, data that would be valuable for primary care physicians. However, no standardized mechanism exists to share this data back to EHR systems. Standardized FHIR profiles for digital therapeutic data would enable seamless information exchange,	startups and academic research groups, democratizing AI development. Comprehensive data interoperability would improve AI model quality and generalizability by enabling training on diverse, complete datasets. Current AI models often suffer from narrow training data, leading to algorithms performing well in development but failing when deployed in different contexts. Federated learning infrastructure would enable breakthrough research current data silos prevent. Developing AI for rare diseases requires aggregating patient data across many institutions, but institutional barriers typically make this impossible. Federated learning would enable collaborative model development while respecting data governance, potentially accelerating rare disease research by 5-10 years. Benchmark datasets would accelerate the field by enabling transparent performance comparison.

RFI Question	GCI Recommendation	Supporting Evidence	Expected Impact
		allowing AI-generated insights from specialized applications to inform broader clinical decision-making.	
Question 9: **What challenges within health care do patients and caregivers wish to see addressed by the adoption and use of AI in clinical care? Equally, what concerns do patients and caregivers have related to the adoption and use of AI in clinical care?**	Patients and caregivers identify access, convenience, and communication as healthcare challenges they most want AI to address, while harboring concerns about privacy, algorithmic bias, transparency, and accountability. **Patient Priorities:** Access barriers (long waits, geographic distances); convenience for chronic condition management (simplifying medication management, coordinating care across providers); communication challenges (difficulty reaching providers between visits, confusion about instructions). Patients want AI to address the fundamental human experience of healthcare: feeling heard, respected, and involved in decision-making. Patients see potential for AI to provide more personalized care by analyzing their individual data and preferences rather than applying inflexible protocols. **Patient Concerns:** Privacy and data security (who has access to	Behavioral health platforms regularly survey patients about experience with AI-enabled features like automated screening, therapist matching algorithms, and progress tracking. Patients overwhelmingly appreciate AI features reducing barriers to accessing mental health care. However, platforms have found patients want transparency about when AI is making decisions versus human clinicians. When a therapist matching algorithm is clearly explained and patients understand they can request human override, acceptance is high. AI-enabled physical therapy platforms reveal patient enthusiasm for AI that enhances convenience while maintaining quality. Patients using AI-powered home therapy consistently report satisfaction with flexibility of exercising on their own schedules and appreciate real-time AI feedback on exercise form. However, patient acceptance increased dramatically when messaging emphasized AI augments rather than replacing human physical therapists, and when easy access to human clinicians was added.	For AI tools deployed in patient-facing applications, HHS should require clear disclosure about AI involvement, written in plain language patients can understand. Creating patient-friendly educational resources about healthcare AI would build health literacy and reduce anxiety about AI-mediated care. HHS also should ensure HIPAA regulations clearly apply to AI systems and enforcement includes algorithmic systems, not just traditional data storage, and transmission. HHS should require AI tools used in clinical care undergo fairness testing across demographic subgroups before deployment and continuous monitoring post-deployment, with public reporting of performance disparities. HHS should clarify liability frameworks so patients understand their rights when AI is involved in their care and ensure medical malpractice systems can appropriately assign responsibility across AI developers, healthcare organizations, and individual clinicians. Most fundamentally, HHS should ensure AI adoption genuinely serves patient interests rather than primarily benefiting healthcare organizations or technology vendors.

RFI Question	GCI Recommendation	Supporting Evidence	Expected Impact
	health information, whether AI systems might expose sensitive data); algorithmic bias (AI systems might provide inferior care to certain groups); transparency and explainability (understanding why AI recommended a specific treatment); accountability (if something goes wrong, who is responsible?).	Patient-facing AI agent research found patient acceptance depends heavily on appropriate use cases. Patients are comfortable with AI handling routine administrative tasks but want human clinicians involved in complex medical decisions, emotionally sensitive conversations, and situations involving uncertainty or judgment. Transparency about AI limitations proves crucial; when AI systems explicitly acknowledge what they cannot do and when human involvement is needed, patient trust increases.	
Question 10: **Are there specific areas of AI research that HHS should prioritize to accelerate the adoption of AI as part of clinical care?** **a. Are there published findings about the impact of adopted AI**	HHS should prioritize three critical research areas representing current bottlenecks to safe, effective AI adoption: **(1) Algorithmic Fairness and Bias Mitigation:** Robust methodologies for detecting, measuring, and mitigating bias in healthcare AI remain underdeveloped. Research should create continuous monitoring approaches detecting fairness degradation as AI systems are deployed across care settings. **(2) Human-AI Collaboration and Clinical Decision-Making:** Most AI research neglects human	Radiology AI companies have invested substantially in fairness research after discovering models showed performance disparities across different CT scanner manufacturers and imaging protocols, creating inequities in care quality. Industry-wide research on fairness testing for medical imaging AI would enable all developers to deploy more equitable systems. Research on optimal human-AI collaboration would help radiology departments train staff more effectively. Generative AI work for healthcare has revealed how little existing research addresses safety concerns specific to language models in clinical contexts. Clinical AI requires additional safety considerations: ensuring AI doesn't	Prioritizing these research areas would accelerate AI adoption by addressing most significant technical barriers to safe, equitable deployment. Investment in fairness and bias mitigation research would produce standardized methodologies AI developers could apply systematically, reducing risk of deploying discriminatory algorithms. The expected outcome is that AI systems demonstrably provide equitable care across all patient populations. Human-AI collaboration research would improve actual clinical utility of AI tools by ensuring they augment rather than impede clinician judgment. Expected impact is higher adoption rates as clinicians develop confidence in AI tools and improved patient outcomes from more effective human-AI collaboration.

RFI Question	GCI Recommendation	Supporting Evidence	Expected Impact
tools and their use in clinical care? **b. How does the literature approach the costs, benefits, and transfers of using AI as part of clinical care?**	factors that determine whether AI tools improve clinical care. Research should develop frameworks for shared decision-making between clinicians, patients, and AI systems. **(3) AI Safety and Robustness Under Distribution Shift:** Healthcare data distributions shift constantly as patient populations change and clinical practices evolve. Research should develop methods for detecting when AI model performance degrades due to distribution shift and techniques for adapting models to new distributions.	provide medical advice beyond its competence, appropriately expresses uncertainty, escalates to human clinicians when needed, and communicates in ways patients with varying health literacy can understand. The entire industry would benefit from shared research frameworks and publicly available benchmarks for evaluating clinical safety of generative AI systems. Healthcare infrastructure work has demonstrated the challenge of distribution shift. AI models sometimes perform inconsistently when deployed across different healthcare settings due to differences in how data is recorded and how clinical workflows operate. Research on automated methods for continuously monitoring AI performance would enable more robust AI deployments across the United States.	AI safety and robustness research would enable confident deployment of AI across diverse healthcare settings by providing methods to detect and address performance degradation before patient harm occurs. The expected outcome is more reliable AI systems maintaining performance across diverse populations and care settings.

Conclusion

The full alignment across HHS agencies through OneHHS presents a historic opportunity to transform U.S. healthcare through AI innovation, and the tools and talent to do so already exist. By establishing tiered regulatory frameworks, expanding value-based payment models, mandating interoperability, supporting industry-driven certification, and investing in critical research, HHS can foster a competitive marketplace that accelerates innovation while protecting patient interests and improving outcomes across all demographics and care settings.

Our message is simple: the tools exist. The talent exists. The policy pathways must now align to ensure these assets can scale across every level of the healthcare system.

The General Catalyst Institute and our ecosystem of portfolio companies stand ready to partner with HHS in this work, bringing real-world data from tens of millions of patients and a track record of health system partnerships to support pilots, testbeds, and ongoing feedback loops as policy implementation advances.

Together, HHS and the private sector can shift the U.S. healthcare system toward **clinical abundance**, ensuring the tools of modern medicine are distributed fairly, used safely, and scaled effectively for every patient. We believe the moment is now and that partnership is the way to make U.S. healthcare truly work for all.

Respectfully submitted,

GENERAL CATALYST INSTITUTE

The General Catalyst Institute
800 Connective Ave NW, Suite 300
Washington, DC 20006
gcinstitute@generalcatalyst.com

Table of Acronyms

Acronym	Definition
ACO	Accountable Care Organization
ADT	Admission, Discharge, and Transfer
AG	Attorney General
AI	Artificial Intelligence
AHRQ	Agency for Healthcare Research and Quality
API	Application Programming Interface
APM	Alternative Payment Model
ASTP	Assistant Secretary for Technology Policy
CFO	Chief Financial Officer
CFR	Code of Federal Regulations
CMIO	Chief Medical Information Officer
CIO	Chief Information Officer
CMS	Centers for Medicare & Medicaid Services
CPT	Current Procedural Terminology
CRADA	Cooperative Research and Development Agreement
CT	Computed Tomography
DEA	Drug Enforcement Administration
DPC	Data at the Point of Care
ED	Emergency Department
EHI	Electronic Health Information
EHR	Electronic Health Record
EMR	Electronic Medical Record
FDA	Food and Drug Administration

FHIR	Fast Healthcare Interoperability Resources
GCI	General Catalyst Institute
HATCo	Health Assurance Transformation Company
HHS	Department of Health and Human Services
HIE	Health Information Exchange
HIPAA	Health Insurance Portability and Accountability Act
HL7	Health Level Seven International
IT	Information Technology
LCD	Local Coverage Determination
MSSP	Medicare Shared Savings Program
NCD	National Coverage Determination
NIH	National Institutes of Health
NIST	National Institute of Standards and Technology
NQF	National Quality Forum
NPPES	National Plan and Provider Enumeration System
OMB	Office of Management and Budget
ONC	Office of the National Coordinator for Health Information Technology
PACS	Picture Archiving and Communication System
R&D	Research and Development
RFI	Request for Information
RFP	Request for Proposal
SDOH	Social Determinants of Health
SNAP	Supplemental Nutrition Assistance Program
USCDI	United States Core Data for Interoperability
WIC	Women, Infants, and Children (Program)