Humanoid AIs in the Research Library: Robots, Research Methods, and the Humanities in the 21st Century

Schmidt Sciences Humanities AI Capabilities Grant • Level I • $250,000 • 24 Months University of California, Riverside

PI: Raymond Uzwyshyn, Ph.D., MBA, MLIS **Co-PIs:** Phoenix Alexander, Ph.D. • Barbara Martinez Neda, Ph.D., M.Sc. • Brendon Wheeler, M Sc. (USC) • Alvaro Alvarez (MBA MSc)

Senior Personnel: Inyoung Shin, Ph.D., M.Sc. (Yale). • Jing Han, Ph.D.

---

1. Problem

Science Fiction writers spend their careers asking whether machines can think, feel, and dream. Many of their manuscripts sit in University of California, Riverside Eaton Collection of Science Fiction and Fantasy — the world's largest publicly accessible collection of speculative literature, over 300,000 items spanning a century of sustained meditation on artificial life. In 2026 though, many of their questions are no longer speculative but raise large questions for the humanities. AI-enhanced humanoid robots walk, talk, sometimes speaking twenty or more languages easily and can recognize faces and emotions, grasp objects, and even be taught to navigate complex environments. The 2026 Unitree G1, with 43 degrees of freedom and onboard NVIDIA AI processing also costs less than a mid-range compact vehicle.

These machines though were built by engineers, for engineers: warehouse picking, industrial inspection and with little thought for university or research environments. All of the major frontier LLM's though, have 'research' functions built out and pass Ph.D. benchmarks on multiple levels. The humanities and social sciences are disciplines that study language, narrative, material culture, embodied knowledge, and the full multimodal texture of human experience. That could contribute majorly to a variety of interdisciplinary AI/robotics areas. Up to now, they have had little opportunity to shape or examine what AI enhanced humanoid robots may accomplish. This includes chance to discover what they might do for scholarship, learning or research in a humanities interdisciplinary research environment and academic library.

This project brings AI-enhanced humanoid robots into a tier 1 American Research library (ARL) and Carnegie 1 Research University together for the first time. The research sets up and test five AI enhanced robots against various structured scholarly tasks drawn from UCR Libraries' collections and archives. It measures results against performance of trained human staff and students accomplishing identical tasks and working together. Where robots succeed, we will document new methods for humanities research and human/ai collaboration. Where they fail, we will produce benchmarks that tell AI developers what scholarship and research and human/AI

robotics collaboration demands. The library is already the most interdisciplinary space on campus. In this project it becomes both research laboratory for and proving ground for research enhanced technology methods for the second quarter of the 21st century

HAVI Connection. HAVI funds new AI tools for humanities research. We test whether the newest class of AI enhanced form factor and platform — the humanoid robot — can serve humanities and social science scholars and the wider university in an interdisciplinary context and space. The RFP's own example captures our approach: "can we develop a new kind of set of research methodologies that uses AI enhanced robots to help our researchers work and students learn?"

Humanities Research Question. How can AI-enhanced humanoid robots serve as research collaborators with humans within a 21st-century research library — navigating archives, assisting oral history collections, documenting material culture, embodying research librarian or author personas, and partnering with subject specialists and researchers — and what domain knowledge must humanities scholars contribute to make these machines scholarly?

AI Research Question. Can a humanities-derived human/ai collaboration benchmark for embodied AI robotics reveal affordances, possibilities but also deficits in multimodal understanding that engineering benchmarks miss, and can datasets from structured humanities researchers and tasks improve humanoid AI through fine-tuning and deeper review and examination of the new medium specificity allowed by the AI enhanced humanoid robot form factor?

Section 2: Approach and Methods

Five AI-enhanced robotic platforms. Four research tasks distributed across two libraries and three labs. A biweekly research colloquium with separate team working sessions. Summer research intensive where graduate students and faculty build disciplinary projects with humanoid AI. Two years. Biweekly progress reports throughout.

2.1 UCR Libraries: Orbach, Rivera, The Robotics Laboratory and Makerspace

This project distributes humanoid robots across UCR Libraries' two main facilities and three specialized labs, embedding them in the daily practice of research, reference, and discovery.

Tomás Rivera Library houses Special Collections and University Archives on the fourth floor — including the Eaton Collection of Science Fiction and Fantasy, the world's largest publicly accessible collection of speculative literature (300,000+ items); the Tuskegee Airmen oral history collection; Tomás Rivera papers; Rupert Coston Library of the American Indian; and Inland Empire California ‘Citrus’ and Water resource’ archives. Rivera also provides the library's primary public research services across humanities and social sciences.

Orbach Science Library houses the Creat'R Lab makerspace, Robotics Lab, 3DXP Lab (virtual reality and 3D scanning), STAR Lab (high-performance and AI computing), and the Geospatial Resources Collection. Orbach serves STEM and interdisciplinary researchers and provides the technical infrastructure for platform configuration and testing.

Both libraries connect to UCR's digital research infrastructure: UC Library Search, the Calisphere digital repository, eScholarship, and the Online Archive of California. Robots operating in physical library spaces can query these systems simultaneously, bridging material and digital scholarship in real time.

2.2 Five Robotic Platforms, Distributed

#	Platform	Location	Role	Cost
1	**Unitree G1 EDU Standard**	Orbach: Robotics Lab + public spaces	Flagship humanoid. Research collaboration. Makerspace integration. Science reference.	$43,500
2	**NAO v6**	Rivera: Eaton Collection + Special Collections	Author persona tours. Multilingual dialogue. Archival engagement.	$10,000
3	**Unitree R1 EDU**	Rivera: Reference + public spaces	Humanities/social science reference partnership. Patron engagement.	$12,000
4	**Misty II**	Rivera or Orbach (mobile deployment)	Customizable community interaction. Oral history assistance. Flexible assignment across tasks.	$3,500
5	**Rover Mini (ROS2)**	Orbach: Creat'R Lab + stacks	Indoor autonomous documentation. Collection survey. Wander Rover research continuity.	$3,200
6	**Reserve**	TBD	Next-gen humanoid or additional peripherals as needs emerge.	$8,000

Total: ~$80,200. Platforms distributed across both libraries. Each robot has a home location but moves between sites as tasks require. All connect to LLMs via Python/ROS2 APIs.

Peripherals/equipment buffer ($8,000): microphones, cameras, charging stations, compute modules, mounts, cables, replacement parts, unanticipated needs.

2.3 Four Research Tasks

Task A. Eaton Collection: Author Persona and Archival Tour.

A NAO robot, connected to a language model fine-tuned on key science fiction archives papers and the Eaton Collection's catalog, guides visitors through the collection as a luminary science fiction author persona (Phlip K. Dick, Octavia Butler, Ursula Leguin) — narrating manuscripts,

cross-referencing related items, and responding to questions. This task tests persona embodiment, multimodal narrative, and the integration of physical navigation with digital collection metadata. One task, but it is the signature demonstration — science fiction's most famous exploration of artificial consciousness, brought to life through the archive that preserves it.

Baseline: trained archivist conducting a standard guided tour and fine tuning**.**

Task B. Rivera Library: Humanities and Social Science Research Reference.

A Unitree R1 EDU and Misty II work alongside human subject librarians at Rivera and Orbach Library's research reference services — answering research questions, directing patrons to relevant collections and UCR’s 600 database collection, providing subject-specific guidance across humanities and social science disciplines, and querying UC Library Search and eScholarship in real time. This task tests whether humanoid robots can collaborate, replace, help or augment the research librarian's capacity to serve students and faculty — the core function of any research library.

Baseline: subject librarian working without robotic assistance.

Task C. Orbach Science Library: Interdisciplinary Research Collaboration and Material Documentation.

The Unitree G1 — the project's flagship humanoid, with full-scale human height, 43 degrees of freedom, LiDAR, and onboard NVIDIA AI — operates in the Orbach Robotics Lab and Creat'R Lab makerspace and extends into Orbach's public research spaces. This task tests the G1's capacity for hands-on research collaboration: 3D object documentation integrating the Creat'R Lab's scanning and fabrication tools; physical handling of materials; and interdisciplinary research support for Humanities, STEM and social science patrons. The Rover Mini conducts autonomous indoor documentation and collection survey.

Baseline: trained specialist using standard documentation methods**.**

Task D. AI Integration: Fine-Tuning, Benchmarks, and Dataset Construction.

This cross-cutting task spans all platforms and locations. Various technical leads develop the technical AI work that makes the other three tasks possible and produces the project's primary AI outputs: fine-tuning large language models on Eaton Collection materials and UCR's digital collections; connecting fine-tuned models to each robot's speech, gesture, and navigation systems; constructing the Humanities Research Instrument Benchmark (HRIB) from structured data across all tasks; and curating the Multimodal Humanities Interaction Dataset (MHID) with domain-expert annotations. This is the engine that converts humanities domain knowledge into measurable AI advancement.

Baseline: standard off-the-shelf robot configurations without humanities-informed fine-tuning.

2.4 Methodology: Two-Year Plan

Year 1, Begin September 2026: Build the Foundation.

Phase 1 (Months 1–4): Acquire and configure five platforms. Deploy NAO and R1 EDU at Rivera Library. Deploy G1 and Rover Mini at Orbach. Configure Misty II for flexible assignment. Core staff — Wheeler, Alvarez, Coordinator, GRA, and Creat'R Lab student workers — handle integration. Each co-PI designs a research protocol for their task. Martinez Neda begins fine-tuning LLMs on Eaton Collection materials. Alexander develops the Science Fiction persona script and archival tour structure. IRB submitted Month 1.

Phase 2 (Months 5–8): Trained human specialists perform each task using standard non-robotic methods, establishing baselines for comparison. Each platform undergoes pilot testing in its home location. Protocols refined based on pilot data.

Phase 3 (Months 9–12): Continued testing, refinement, and data collection. Technical leads complete first-generation fine-tuned models and connects them to robot APIs. Competitive call for Year 2 Summer Research Intensive participants issued Month 10. Biweekly progress report: platform status, baseline data, protocol revisions, model performance.

Year 2, September 2027: Test, Collaborate, Publish.

Phase 4 (Months 13–15): **Biweekly Research Colloquium** (Winter Quarter). Twenty to thirty participants selected by competitive application — graduate students and faculty from humanities, social sciences, and engineering. Full group meets monthly (every other week) for structured presentations, cross-team review, and benchmark data discussion. Between full-group meetings, disciplinary teams meet bi-weekly with their assigned robots to advance their specific work.

Phase 5 (Months 16–20): **Summer Research Intensive** (Spring/Summer). The core production period. Disciplinary teams — each composed of graduate students and faculty from a specific humanities or social science field — work intensively with assigned robotic platforms to build research projects within their own disciplines. An anthropology team uses a humanoid for ethnographic data collection. A media studies team explores robot-mediated narrative. A history team tests archival navigation methods. A business team examines AI-assisted organizational research. A research library team completes a library project. Each team produces a working paper documenting methods, findings, and benchmark data from their disciplinary perspective. The PI and co-PIs provide mentorship and ensure cross-pollination across teams. Weekly team meetings; biweekly full-group sessions continue.

Phase 6 (Months 20–24): Formalize the **Humanities Research Instrument Benchmark (HRIB)** with human baseline comparisons across all four tasks. Curate the **Multimodal Humanities Interaction Dataset (MHID)** with domain-expert annotations. Final model evaluation. Release all outputs open-access via GitHub and Hugging Face. Submit 3–5 papers. Principals given opportunity to present at 2–3 conferences. Prepare Level II proposal for UC-wide expansion.

Meeting Structure:

- Year 1: Core team (PI + co-PIs + Coordinator + GRA) meet monthly with progress reports.
- Year 2 Winter: Full colloquium meets monthly. Disciplinary teams meet separately on alternating weeks with their robots.
- Year 2 Spring/Summer: Disciplinary summer camp group meet twice weekly with robots. Full group meets biweekly for cross-team review and integration.
- Year 2 Fall: Publication and conference focus. Monthly full-group meetings. Final Project Report

---

Two Year Timeline

Months	Activities	Outputs
Year 1 1–4	Platform acquisition. Deployment: NAO + R1 at Rivera, G1 + Rover at Orbach. IRB. Protocol design. LLM fine-tuning investigation and beginning.	5 platforms configured and placed. IRB approved. 4 protocols drafted. Science Fiction persona v1.
5–8	Human baselines collected. Pilot testing at both libraries. Protocol refinement	Baselines for all 4 tasks. Pilot data. Refined protocols.
9–12	Continued testing. Fine-tuned models connected to robots. Colloquium call issued. Participants selected.	Models v1 deployed. Competitive call completed. 20–30 participants selected.
Year 2 13–15	MONTHLY RESEARCH COLLOQUIUM (Winter). Full Biweekly team sessions between.	Cross-team benchmark discussions. Disciplinary teams formed. Preliminary data.
16–20	SUMMER RESEARCH INTENSIVE. Humanities Disciplinary teams build projects with robots. Weekly team meetings, biweekly full group.	Working papers from each team. Benchmark data across disciplines. Public presentation.
20–24	HRIB formalized. MHID curated. Papers submitted. Conference presentations. Level II prep.	HRIB + MHID open-access. 4–6 papers. All 5 principals at conferences. Level II drafted.

Midpoint (Month 12): Five platforms operational across both libraries. Human baselines complete. Fine-tuned models connected. Colloquium participants selected.

Endpoint (Month 24): HRIB and MHID released open-access. 3–5 papers submitted. Principals present at AI & Robotics related conferences. 20–30 researchers trained. Level II proposal ready for UC-wide expansion.

---

3. Environmental Scan

AI enhanced Humanoid robotics reached commercial viability in 2025–2026. Unitree shipped over 5,000 G1 units and plans 10,000–20,000 in 2026. Over two dozen models from Chinese and American manufacturers are now available. Yet deployment remains almost entirely industrial: Agility Digit in Amazon warehouses, Figure 02 at BMW. No research library or university wider public community center has integrated humanoid robots. No humanities program has tested robots against scholarly standards for research help even though Deep Research AI models have been around for two years. No benchmark exists for evaluating embodied AI on humanistic tasks.

HAVI's 2025 portfolio defines the closest precedents. Casona et al. use AI remote sensing as a field instrument for archaeological discovery; we extend field instrumentation from satellites to ground-level humanoid platforms. Underwood and Wilkens benchmark language models against humanities standards for cultural and historical reasoning; this project extends humanities benchmarking from text-based AI to embodied robotics AI. Casey et al. build AI tools for multimodal newspaper archives; we work with collections that help human researchers in physical spaces that no text-based AI can navigate.

What is new: (1) Humanities-derived research human/robot collaboration benchmark for embodied AI. (2) First research library R&D integration of humanoid robots. (3) First robotic-human collaboration/comparison for scholarly research tasks. (4) The world's foremost science fiction collection as special collection test corpus — closing the circle between speculative imagination and material reality. (5) A new generation of humanities and social science researchers beginning to work out general models for knowledge work and humanities research with humanoid AI.

---

4. Impact

On the Humanities and Social Sciences. This project models a 21st-century research library transformed by humanoid intelligence. A robot channels science fiction personas to help researchers find collections. A full-scale humanoid AI assistant works beside a subject librarian at the reference desk, querying digital databases while the librarian handles the human conversation. These are structured experiments producing benchmarks, datasets, and publications. The HRIB benchmark enables any institution to evaluate humanoid platforms for scholarly use and affordances. Twenty to thirty participants become the first humanities and social science researchers trained for this emerging paradigm — a new understanding of a new

medium specificity: AI humanoid robotics. The project generates genuine research findings and three to five publications and open instrument specifications extend the work beyond UCR.

On Artificial Intelligence. The HRIB benchmark exposes where embodied AI succeeds or fails on tasks requiring interpretive ambiguity, gesture-speech integration, multilingual processing, persona performance, and contextual reasoning. The MHID dataset, annotated by domain experts, provides training data and fine-tuning experiments give lists for best practices for humanoid/researcher collaboration performance. All benchmarks, datasets, and models released via Arxiv, GitHub, Hugging Face and Data Dryad.

---

5. Team

Raymond Uzwyshyn, Ph.D., MBA, MLIS — PI. Acting AUL for Research & Technology / Director of Research Services, UCR. Ph.D. Media Studies (NYU) doctoral training in media theory, history of technology; 25 years, IT Project Management leadership (MBA) at Academic R1/ARL institutions; directs UCR's AI, robotics, GIS, and makerspace facilities. Co-editor, *New Horizons in AI for Libraries* (IFLA/De Gruyter, 2025). Amazon Think Big Science grant, autonomous rover research (2026). **Role:** Project direction; benchmark, proposal.

Barbara Martinez Neda, Ph.D., M.Sc. — Co-PI. Research Data Scientist, UCR. Ph.D. Machine Learning. M.Sc. Computer Engineering. Expertise in large language models, autonomous systems, Python. **Role:** AI model fine-tuning; LLM-robot integration; dataset construction; quantitative benchmark analysis..

Phoenix Alexander, Ph.D. — Co-PI. Jay Kay and Doris Klein Librarian for Science Fiction and Fantasy, UCR. Eaton Collection Curator. Published SF author (*F&SF*, *Lightspeed*, *Escape Pod*). SFWA History Committee. Presenter at Worldcon and San Diego Comic-Con. **Role:** Task A (archival guide and persona embodiment); Eaton Collection research design; public engagement. 5% effort.

Brendon Wheeler — Co-PI. Maker Services Coordinator, UCR Creat'R Lab. CAD, 3D printing, embedded systems, Python/C++. Led UCR Robotics Camps 2023–2025 training 40+ students in AI and autonomous systems. **Role:** Task C (3D documentation); platform configuration, integration, and maintenance; lab operations. Supervises Creat'R Lab student workers who support daily robot operations.

Inyoung Shin, Ph.D., M.S. — Senior Personnel. AI Research Librarian, UCR. Ph.D. Communication (Rutgers); M.S. Computer Science (Yale). Specialization: human–robot interaction, new technologies. **Role:** HRI methodology advisor; benchmark evaluation protocol design. Contributes from existing position.

Alvaro Alvarez — Co-PI. Innovative Media Librarian, UCR. Interactive media production, GIS-based learning modules, NSF STEM outreach experience. **Role:** Multimedia data capture; 3DXP Lab integration, Robotics Integration; public demonstrations; colloquium media

documentation.

Jing Han, Ph.D. — Senior Personnel. Digital Scholarship Librarian, UCR. Leads UCR's Digital Scholarship Program across humanities, social sciences, and STEM. Computational text analysis, data visualization, digital media production. Hosts *The Reference Desk* podcast. **Role:** Digital scholarship integration advisor; dataset curation support; eScholarship and Calisphere connectivity. Contributes from existing position.

Project Coordinator / HCI Engineer (new hire, 0.5 FTE) **and Humanities student workers (internal**). Manages daily technical integration across five platforms, colloquium logistics, biweekly deliverable tracking, and data infrastructure. Required: HCI or robotics engineering background; ROS2 familiarity. Works closely with Wheeler and Alvarez on platform operations. Humanities Robotics lab student workers will also be empowered by learning how to deploy, interact and benchmark robotics systems embedded with artificial intelligence allowing them to understand how systems are applied for later workforce development, on ramps gaining valuable technical expertise.

Graduate Research Assistant (1, both years). Humanities or social science background preferred. Full colloquium participation. Contributes to benchmarks, annotations, and publications, humanities faculty and staff relations. Weekly mentorship from PI. Supported by existing Creat'R Lab student workers for daily robot operations.

Colloquium Participants (20–30). Selected by competitive application from UCR graduate humanities students and research humanities faculty. Balanced across humanities (media studies, history, anthropology, art history, English etc.), social sciences (sociology, business, public policy), and engineering (CS, EE). Faculty serve as domain consultants. Participants produce white papers and gain direct experience with humanoid AI — the first cohort of humanities and social science researchers prepared for this emerging research paradigm.

Mentorship and Collaboration. The PI has mentored students across disciplines throughout a career spanning six R1 institutions. The summer camp robotics structure ensures genuine interdisciplinary collaboration: humanities scholars define what the robots must achieve; technologists build and test; progress reports ensure domain knowledge shapes development at every step. Student success and workforce development are central to the project's design.

---

6. Outreach and Dissemination

Final Products.

Humanities Research Instrument Benchmark (HRIB) — open-access, eScholarship, Data Dryad.

Multimodal Humanities Interaction Dataset (MHID) — Hugging Face.Data Dryad, Escholarship

Fine-tuned AI models with documented performance data.

Platform specification documents for each of five robots. Arxiv

4–6 peer-reviewed publications targeting *Digital Humanities Quarterly*, *Digital Scholarship in the Humanities*, *ACM/IEEE HRI Conference*, *International Journal of Social Robotics*. NeurIPS

Project website with video documentation of all four tasks.

Conference Presentations. All principal investigators present results at 1–2 conferences by Year 2: ACM Conference on Human-Robot Interaction; ADHO Digital Humanities; HAVI community events; IFLA World Library and Information Congress. Budget of $25,000 supports full team participation.

Broader Reach. UC system workshops via the PI's existing UCSF Generative AI network and UC committee roles. Open-access preprints (arXiv, SocArXiv). IFLA international research library community via the PI's membership in IFLA and editorship of *Trends and Issues in Library Technology*. Eaton Collection public events organized by Alexander, including the biennial Eaton Conference on Speculative Fiction. All benchmarks, datasets, and specifications documented for adoption by other research libraries.

References

Bartneck, C., Kulić, D., Croft, E., & Zoghbi, S. (2009). Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. *International Journal of Social Robotics*, 1(1), 71–81.

Burgard, W. et al. (1999). Experiences with an interactive museum tour-guide robot. *Artificial Intelligence*, 114(1–2), 3–55.

Casana, J. et al. (2025). Discovering global archaeological heritage using AI and remote sensing. Schmidt Sciences HAVI.

Casey, J. et al. (2025). Communities in the loop: AI for cultures and contexts in multimodal archives. Schmidt Sciences HAVI.

Drucker, J. (2021). *The Digital Humanities Coursebook*. Routledge.

Elkins, K. et al. (2025). Archival intelligence: Rescuing New Orleans' endangered cultural legacy. Schmidt Sciences HAVI.

Faber, F. & Bennewitz, M. (2009). The humanoid museum tour guide Robotinho. *Proc. IEEE RO-MAN*.

Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press.

Hayles, N. K. (2012). *How We Think: Digital Media and Contemporary Technogenesis*. University of Chicago Press.

Kittler, F. (1999). *Gramophone, Film, Typewriter*. Stanford University Press.

Latour, B. & Woolgar, S. (1986). *Laboratory Life: The Construction of Scientific Facts*. 2nd ed. Princeton University Press.

Matarić, M. J. (2007). *The Robotics Primer*. MIT Press.

Suchman, L. (2007). *Human-Machine Reconfigurations: Plans and Situated Actions*. 2nd ed. Cambridge University Press.

Thrun, S. et al. (2000). Probabilistic algorithms and the interactive museum tour-guide robot Minerva. *International Journal of Robotics Research*, 19(11), 972–999.

Underwood, T. & Wilkens, M. (2025). Artificial intelligence for cultural and historical reasoning. Schmidt Sciences HAVI.

Uzwyshyn, R. et al. (Eds.). (2025). *New Horizons in AI for Libraries*. IFLA/De Gruyter.

Wang, X. et al. (2025). Social robots for child development. *Humanities and Social Sciences Communications*, 12, 1411.

Zlotowski, J. et al. (2015). Anthropomorphism: Opportunities and challenges in HRI. *International Journal of Social Robotics*, 7(3), 347–360.

Schmidt Grant Form Fields

Level: Level I

Budget: $250,000

Start: 09/26

End: 08/28

Keywords: humanoid robots, research library, multimodal benchmark, speculative literature, embodied AI

Disciplines: media studies, computer science, machine learning, speculative fiction studies, digital humanities

Abstract

Science Fiction writers imagine humanoid robots. Their papers sit in UCR's Eaton Collection, the world's largest publicly accessible collection of speculative literature. In 2026, those imagined machines are real. This project integrates five AI-enhanced humanoid robotic platforms across UCR Libraries and tests them against four structured scholarly tasks drawn from the Eaton Collection and Special Collections: humanoid navigation with author persona embodiment and research reference partnership. Each task benchmarks robotic against trained human performance. A ten-week Spring Research Colloquium brings selected graduate students and faculty from humanities, social sciences, and engineering to test and refine these platforms and summer camp in the second-year trains researchers working in the humanoids with robots. Outputs: the first Humanities Research Instrument Benchmark for embodied AI, a multimodal dataset, fine-tuned models, and 4-6 peer-reviewed publications.

Humanities Research Question (under 1000 chars):

How can AI-enhanced humanoid robots serve humanities and social science research within a 21st-century research library and university? UCR Libraries houses the world's largest publicly accessible collection of speculative literature alongside 750 archival collections. We test five robotic platforms against four tasks drawn from collections and robotic navigation with author persona embodiment, multilingual tours, material culture documentation, and research reference partnership alongside human subject librarians. Each task benchmarks robotic against human performance. We produce evaluation frameworks, platform specifications, and a model for the research library transformed by humanoid AI, while training the first cohort of humanities and social science researchers equipped to work with these emerging platforms.

AI Research Question (under 1000 chars):

Can a humanities-derived human/robot collaboration benchmarks for embodied AI reveal affordances and deficits in humanoid robots' multimodal understanding? Humanities research demands interpretive ambiguity, gesture-speech coordination, multilingual processing, persona embodiment, and contextual reasoning. This research constructs human/ai/robotics e Humanities Research Instrument Benchmarks testing five humanoid robots across four scholarly tasks, comparing robotic against and with human alone human/robot and robot only baselines. We produce a multimodal dataset with domain-expert annotations capturing interpretive context absent from robotics training data. We test whether fine-tuning on humanities-informed data improves embodied AI multimodal performance for the next level of development of humanities research and robotics.

Compute: ~400 GPU-hours (H100): LLM fine-tuning on Eaton materials (~250), benchmark evaluation (~100), multimodal model training (~50). Storage: ~2 TB.