GOOD
SYSTEMS

Fiscal Year 2025
Annual Report

The University of Texas at Austin
Research Development
*Vice President for Research, Scholarship
and Creative Endeavors*

Table of Contents

Advancing ethical, trustworthy innovation in AI is our grand challenge.

Artificial intelligence is transforming nearly every industry, academic field and government sector, offering powerful new tools to tackle problems and address complicated issues. While AI promises tremendous benefits, it also poses risks that aren't always apparent. We must look beyond technical capabilities to consider ethical boundaries and social implications.

At The University of Texas at Austin, the Good Systems Grand Challenge is driving efforts to align AI with the public good. Six core research project teams — spanning STEM fields, humanities, social sciences and more — are working to understand how AI is reshaping society and to ensure safe and ethical use as technologies evolve. Teams partner with city and state agencies, industry, nonprofits and community members to define, evaluate and build AI systems that serve the public good and contribute to a global conversation.

Central to this work is a shared commitment to building technologies that foster accountability and consider human integrity. Good Systems defines a "good" system as a human-AI partnership that enhances rather than undermines agency, fosters trust and transparency, and supports the foundations of an open and participatory society.

Good Systems
Highlights in Numbers

EXPANDING NETWORKS

147 active researchers

28 UT departments & disciplines

14 schools, colleges, and units

41 partners

15 hosted events in the past year

ENGAGING STUDENTS

75 undergraduate & graduate student researchers

SCHOLARLY OUTPUT & PUBLICITY

47 scholarly works published in the past year*

20 news articles in the past year

BUILDING CAPACITY

$21M awarded in external funding to date

*acknowledging Good Systems support

Program Achievements

Research

Since launching six core research projects in 2021, Good Systems has been at the forefront of exploring what it means to design ethical AI technologies that benefit society. In Fiscal Year 2025, the six interdisciplinary research teams delivered insights and innovations with wide-ranging applications for developers, corporate leaders, city officials, policymakers, educators, skilled trade workers and people everywhere. The following highlights showcase how UT Austin is helping to shape the future of ethical AI, locally and worldwide.

Being Watched: Embedding Ethics in Public Cameras

As a largely invisible network of surveillance cameras silently observes our daily lives, Good Systems researchers are addressing the ethical challenges posed by the increased use of AI-powered surveillance technologies meant to enhance public safety and optimize municipal resources.

Atlas Wang (Chandra Department of Electrical and Computer Engineering, Cockrell School of Engineering) and his team developed a "differential access model," a framework that restricts who can access surveillance data and for what purposes, designed to customize privacy protections. "Not everyone defines privacy the same way," Wang said. "It varies across cultures, age groups and even personal preferences. Our goal is to make these systems adaptable, so individuals or communities can choose their level of comfort."

Anita Varma (School of Journalism and Media, Moody College of Communications) analyzed public discourse on traffic cameras in California, revealing how these systems disproportionately impact marginalized communities. She proposes an ethical framework for AI development and deployment that centers solidarity by engaging the public in participatory design, addressing the societal impacts of using AI technologies in municipal services.

"Not everyone defines privacy the same way. It varies across cultures, age groups and even personal preferences. Our goal is to make these systems adaptable, so individuals or communities can choose their level of comfort."

— Atlas Wang, Cockrell School of Engineering

Sharon Strover (School of Journalism and Media, Moody College of Communication) was selected by the City of Austin to lead a project developing AI governance and training tools for municipal use. Her team is surveying city employees and designing curriculum to support the ethical use of AI technologies in public services and within its government operations.

Good Systems researchers have developed ethical frameworks for AI-powered surveillance, from customizable privacy models to city-level governance tools and training curricula.

The project team is also liaising with experts from across the US. They convened an in-person meeting with their stakeholder advisory council, which includes university and nonprofit representatives from Cal State Long Beach, UMass Amherst, Data & Society and US Ignite, and shared preliminary findings with GovAI, a national coalition of researchers and civic leaders.

The project team's growing influence in ethical, civic and policy-focused conversations around AI and digital systems was affirmed by invitations to present at several high-impact national and international conferences. A key highlight was inclusion in the American Political Science Association conference, where the team contributed to discussions with scholars, political scientists and other experts on topics such as global AI governance, the use of AI in law enforcement and **AI surveillance**.

READ MORE

Designing Responsible AI Technologies to Protect Information Integrity

Focused on designing, building and testing innovative AI technologies to support journalists, professional fact-checkers and information analysts, this team is developing systems that prioritize fairness and explainability to protect information integrity world-wide.

Researchers built multimodal fact-checkers and visual data analyzers to support journalists, curb misleading information and advance trustworthy, human-centered AI. An analysis of data collected during the 2024 U.S. elections is also underway.

The team made advancements on several tools for identifying and evaluating information accuracy and learned that combining technological innovation with human expertise is key to combating false narratives online. **Dhiraj Murthy** (School of Journalism and Media, Moody College of Communication) developed an mpox dashboard to visualize health-related news and social media information, informed by needs identified by the City of Austin, and co-led the development of **Entendre**, an AI-powered bot detection system designed to identify automated false information networks.

Greg Durrett (Computer Science, College of Natural Sciences) developed a system for fact-checking the outputs of large language models (LLMs) that has the power of ChatGPT-4 at a fraction of the cost. The system, **MiniCheck**, has since been extended and commercialized by Bespoke Labs.

To address the urgent need for tools that promote information integrity in an era of rapidly generated and distributed content, the team expanded their scope to include multimodal fact-checking, enabling analysis of claims embedded in images, charts and graphs. Durrett's partnership with Bespoke Labs also led to the development of Bespoke-MiniChart, a cutting-edge AI Vision-Language Model that can understand, analyze, and answer questions about charts. The model sets a new standard for reading visual data like graphs and infographics, with the potential to create new possibilities for accessibility, education and data reporting.

Josephine "Jo" Lukito (School of Journalism and Media, Moody College of Communication) is expanding knowledge on generative AI and how information is spread across social networks, especially in relation to elections. **Matthew Lease** (School of Information) investigated the multidimensional nature of claim "checkworthiness" and how to better assist fact-checkers given the massive volume of potentially false claims circulating online. He also co-designed an AI-assisted claim-prioritization platform that helps prioritize suspicious posts for human fact-checkers, improving response times and accuracy. As part of this study, Lease co-hosted a series of co-design workshops with professional fact-checkers around the world, to better understand their workflow and generate critical gap analyses that informed the development of these tools.

Researchers received national recognition for advancing the frontiers of responsible AI and language technologies through published papers and presentations at major academic venues. Those include the Conference on Empirical Methods in Natural Language Processing (EMNLP), where Durrett and **Jessy Li**'s (Linguistics, College of Liberal Arts) paper on a new AI tool called QSalience received the 2024 Outstanding Paper Award, and the ACM CHI Conference on Human Factors in Computing Systems, where **Maria De-Arteaga** (Information, Risk and Operation Management, McCombs School of Business) was awarded the Best Paper Honorable Mention.

[READ MORE]

Living and Working with Robots

As robots become more common in homes, workplaces and public spaces, understanding how humans respond to them is essential for designing safe and effective technologies. A multidisciplinary team of UT Austin researchers conducted one of the first studies to combine real-time physiological data with realistic robot encounters to assess how people experience stress during interactions with mobile quadruped (four-legged) robots. The results offer valuable insights for designing safer and more comfortable human-robot interactions.
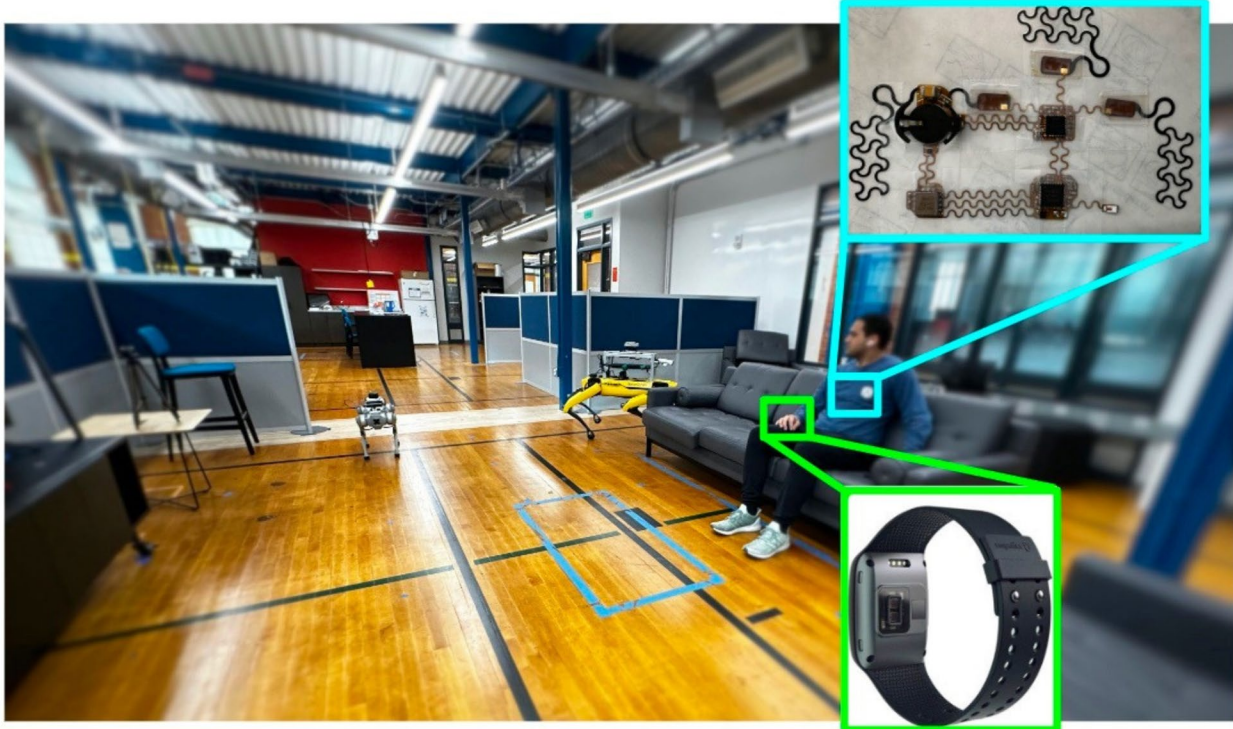
Engineers and communication experts found that brief unpredictability in robot behavior is tolerable, but repeated or prolonged surprises elevate stress, revealing the critical role of behavior design in building safe, emotionally aware robotic systems.

In their study, Nanshu Lu, Luis Sentis and Ryan Gupta (Aerospace Engineering and Engineering Mechanics, Cockrell School of Engineering); Hyonyoung Shin (Chandra Department of Electrical and Computer Engineering, Cockrell School of Engineering); and Keri Stephens and Emily Norman (Communication Studies, Moody College of Communication) measured both physiological stress responses and subjective feelings of stress as participants interacted with or simply coexisted alongside the robots in realistic indoor settings. The team found that participants experienced elevated stress — especially when encountering multiple robots or when the robots were in motion, offering new insights into how robot behavior, movement and presence influence human perceptions of safety.

Findings were shared as a proceedings paper, "**Human stress response and perceived safety during encounters with quadruped robots,**" at the IEEE International Conference on Robot and Human Interactive Communication (RO-MAN), a leading global forum for interdisciplinary research on human-robot interaction. This study lays the groundwork for developing next-generation robots that can adapt to human emotional and physiological cues, advancing the design of technologies that are not only functional but socially and psychologically attuned.

Another project presented at RO-MAN, **"Dobby: A Conversational Service Robot Driven by GPT-4**," highlighted new robotics techniques that enhance conversational ability and task coordination. Dobby is a mobile service robot that can both understand and respond to natural language and perform multi-step tasks, such as guiding visitors through a laboratory while engaging in human-like conversation. The team hopes to deploy the robot as a tour guide at UT's Anna Hiss Gym. The team advanced Dobby's capabilities by integrating large language model architecture with real-time memory, spatial grounding, and "personality expression," enabling it to interact more naturally, recall previous conversations and adapt its behavior based on user engagement during live demonstrations.

Smart Hand Tools: Building the Future of Work with Ethical AI

The future of work isn't about AI replacing human skill; it's about augmenting it. A team of Good Systems researchers is developing smart hand tools equipped with sensing capabilities, IoT (Internet of Things) connectivity and real-time feedback to enhance worker performance, improve safety, boost efficiency and increase accessibility across various sectors. With researchers from the social sciences and engineering, the team is working together to co-develop AI tools with their end users that will help bridge the middle-skill labor gap and contribute to a more secure future of work.

In an era when discussions about automation often center on job displacement, this research offers a counter-narrative, highlighting the nuanced and evolving relationship between humans and artificial intelligence in the workplace. The research study "**Workers Who Care: AI-Enabled Smart Hand Tools in the Skilled Trades**" was published in the Proceedings of the ACM on

This study lays the groundwork for developing next-generation robots that can adapt to human emotional and physiological cues, advancing the design of technologies that are not only functional but socially and psychologically attuned.

Human-Computer Interaction, a leading, peer-reviewed journal that publishes cutting-edge research on how people interact with technology, and will be presented at the ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW) this fall. The study, which examines how expressions of care can inform the design of AI tools, was conducted by **Ken Fleischmann**, **Chelsea McCullough**, **Tina Lassiter** (School of Information) and **Sherri Greenberg** (LBJ School of Public Affairs), in collaboration with the City of Austin. It used qualitative methods, including interviews, focus groups and co-design workshops, to gather insights from city-employed workers and supervisors.

In skilled trades, tools are more than equipment — they reflect pride, shared responsibility and care. Researchers are building AI-enabled smart tools that honor this ethic while supporting safety and learning on the job.

The findings show that everyday gestures of care such as mentoring, proper tool maintenance and clear communication are embedded in the routines of skilled trade workers and essential to workplace safety, efficiency and culture. To support these care-driven practices, the researchers co-developed AI-enhanced "smart" hand tools featuring safety alerts, training simulations and real-time performance feedback that bolster, rather than replace, human expertise.

By involving the actual workers who use these tools, the project emphasizes how human experience should guide the design of AI technology. The results suggest that when AI systems are co-created with end-users,

especially in vital yet underacknowledged fields like the skilled trades, they can support workplace wellbeing, enhance skill development, and lay the groundwork for a future of work that merges technological innovation with stakeholder needs and widely accepted human values.

Led by **Raul Longoria** (Walker Department of Mechanical Engineering, Cockrell School of Engineering), the team also made advancements in smart hand tool design, developing a Prototype Sensing Unit (PSU) that can be placed on any tool, including drills, saws and other rotary tools. This builds on the team's smart rotary power tool design but enables more generalized use by creating sensor units that can be attached to any given tool rather than embedded in just one. The sensor can identify tool types, recognize specific tasks and provide predictive insights that seek to enhance user performance.

As part of this study, four machine learning-based methods were developed to assess user skill in rotary tool use. All were successful at distinguishing novices from experts. However, expert users were harder to rank due to reduced performance variability, a key insight for refining skill-assessment algorithms.

Designing AI to Advance Racial Equity

AI has the potential to either reinforce existing inequalities or help dismantle them; this project works to ensure the latter. By bringing together experts from diverse fields, the team explores how AI could be thoughtfully designed and applied to promote racial equity in key areas such as transportation, public administration and health.

Chandra Bhat (Maseeh Department of Civil, Architectural and Environmental Engineering, Cockrell School of Engineering) and his team focused on pedestrian safety, using AI technologies to better understand how drivers and pedestrians interact at busy intersections. They installed cameras at two locations on campus, analyzed hours of video data, and identified key behavioral patterns to learn how individual and social factors affect pedestrian safety. They also developed a virtual reality (VR) prototype to simulate pedestrian crossing scenarios, offering a way to further investigate pedestrian behavior under controlled conditions. This multi-level approach will allow the team to disentangle the complex dynamic between environmental contexts

In an era when discussions about automation often center on job displacement, this research offers a counter-narrative, highlighting the nuanced and evolving relationship between humans and artificial intelligence in the workplace.

and personal characteristics that shape pedestrian decision-making, providing a nuanced understanding that can inform targeted safety interventions and policy recommendations for diverse demographic groups.

Good Systems researchers are using video and VR to uncover how appearance affects pedestrian safety, as well as the potential risks of coding those patterns into AI. (Image generated using ChatGPT.)

Craig Watkins (School of Journalism and Media, Moody College of Communication) is collaborating with **Kaya de Barbaro** (Psychology, College of Liberal Arts) and her lab to develop algorithms to analyze audio data from wearable sensors that track mother-infant interactions — key indicators of child development. This work builds on previous research on postpartum depression and social determinants of health. To address racial and cultural gaps, such as the underrepresentation of Black mothers in the studies, the team is conducting interviews with healthcare professionals and reviewing literature to design more inclusive research and increase participation among underrepresented groups.

Watkins and his team also produced a health AI strategy paper for the Episcopal Health Foundation, focused on the needs of safety-net healthcare providers — those serving underinsured and underserved populations in

Texas. Based on surveys and interviews with over 200 practitioners, the research identified both the promise and pitfalls of AI in these settings. Key recommendations include building AI literacy among healthcare workers, actively involving rural and underserved communities in the development process, creating robust evaluation metrics and supporting the design of locally relevant AI tools. The goal is to guide ethical, community-informed health AI innovations that strengthen care for Texas's most vulnerable populations.

Min Kyung Lee (School of Information) examined how public sector organizations use technology to support equity goals. Lee and the team **published research** based on interviews with 36 employees from 10 departments in a U.S. city government and found that employees often face challenges in pursuing equity goals and need better tools, data and organizational support. The research highlighted opportunities to make equity a priority when designing government technology used for decision-making. Co-authors on the research included **Angie Zhang** and **Marshanah Taylor** (School of Information); Bhat and **Angela Haddad** (Maseeh Department of Civil, Architectural and Environmental Engineering); and Watkins.

Throughout the year, the team was invited to share their research on national platforms, contributing to important conversations about the role of AI in society. Highlights included **presentations** at the AAAI Conference on Artificial Intelligence, a leading international academic conference in artificial intelligence held annually.

Read more

A Good System for Smart Cities

As cities grow smarter, the question isn't just about what technology can do but whom it benefits. This project seeks to ensure that AI-driven urban development is not only innovative but equitable. Using the City of Austin as a test bed, the team is designing and developing AI tools that help city planners make more informed, inclusive decisions.

A 3D digital twin of Austin, enhanced with real-time traffic data, helps city planners make more informed, inclusive decisions for improving public services such as electricity use and sustainable urban growth.

The team completed a 3D digital twin of Austin and showcased their work at the Annual Meeting of the Transportation Research Board, which is affiliated with the National Academies of Sciences, Engineering, and Medicine and attended by leading transportation researchers. Junfeng Jiao, **Yiming Xu** and **Huihai Wang** (School of Architecture) presented findings on how the digital twin combines traffic, emergency response, noise and other urban data in a detailed 3D environment. Together, these features make the tool a valuable asset for city planners working to improve public services and support sustainable urban growth.

The team also made progress on a large language model designed to support emergency management efforts. Built in partnership with the City of Austin, the model

is being developed for integration into the digital twin, offering a dynamic new tool for crisis response and planning. Early simulations demonstrated the platform's value in planning for road closures and emergency response. A conference paper is in development, and upcoming enhancements will include real-time data feeds and accessibility mapping to further support decision-makers.

The team both expanded the capabilities of OpenCityAI, a chat model trained on public city information, and launched SafeMate, a multimodal AI agent for emergency preparedness. OpenCityAI now supports multilingual feedback (Arabic, Hindi and Vietnamese) and dynamic data visualizations. SafeMate, a voice- and text-based emergency assistant, delivers accurate, real-time safety guidance using FEMA and CDC resources. In benchmark testing, SafeMate outperformed other AI models in accuracy and reliability, and a pilot with local agencies is in development.

The team also organized the third annual Smart Cities and Generative AI Symposium, held in June 2025, bringing together a cross-sector audience of city leaders, AI professionals, researchers, policymakers and community members.

Cross-Cutting Themes

This past year, Good Systems established cross-cutting themes to build bridges across core research projects and identify new lines of inquiry in the field. This year, researchers advanced research under the themes of Knowledge, Generative AI, and Power; Work at the Human-AI Frontier; and Alignment and Safety.

Jared Jensen (School of Journalism and Media, Moody College of Communication) led the Knowledge, Generative AI, and Power theme, focusing on how generative AI is reshaping human creativity, labor and power dynamics. Working primarily with faculty in the Information Integrity, Being Watched and Living and Working with Robots core research projects, Jensen examined the implications of AI-generated content on creative professions, the power dynamics it introduces in platform labor, and how grassroots organizations are responding. He also conducted an in-depth interview study with Good

The team both expanded the capabilities of OpenCityAI, a chat model trained on public city information, and launched SafeMate, a multimodal AI agent for emergency preparedness.

Systems researchers to understand how ethics are communicated and negotiated across interdisciplinary AI development teams. **Early insights from the research,** which will inform how diverse AI development teams can more effectively navigate ethical issues together, were presented at the International Communication Association conference in June.

Lingyuan Li (School of Information) led the Work at the Human-AI Frontier theme by investigating how AI-enabled hand tools are transforming skilled labor. Her research explored how these tools mediate relationships between workers, institutions and data, raising critical questions about surveillance, autonomy and responsible design. By initiating collaborations with the Smart Hand Tools core research project and exploring synergies with the Being Watched, Smart Cities and Living and Working with Robots teams, Li helped identify opportunities for future research and funding that align with Good Systems' mission of ethical AI design.

Brad Knox (Computer Science, College of Natural Sciences) led the Alignment and Safety theme, investigating how to ensure AI systems behave in ways that reflect human values and interests. Knox collaborated with the Living and Working with Robots team to explore the risks of AI companionship, identifying the specific traits of AI companions that could harm users, disrupt their relationships or pose broader societal concerns. With AI companions becoming more prevalent, Knox's research addresses the ethical significance of human-AI interaction. Knox and his colleagues won an outstanding paper award at the Reinforcement Learning Conference in the category of **Emerging Topics in Reinforcement Learning**, and Knox presented research on **LLMs and clarifying questions**, at the Thirteenth International Conference on Learning Representations and the Thirty-Eighth Annual Conference of Neural Information Processing Systems. He was also awarded a grant from Open Philanthropy to conduct foundational research on beliefs, values and decision-making of artificial agents.

Read more

Campus-wide Engagement

2025 Good Systems Symposium: Exploring AI for the Public Good

In April, the fifth annual 2025 Good Systems Symposium brought together more than 180 faculty, researchers, students, professionals and community members to explore the pressing questions at the intersection of artificial intelligence and society. As the signature annual event of UT Austin's Good Systems research grand challenge, the symposium offered a dynamic forum for sharing research, sparking dialogue and building community around the responsible design of AI technologies.

Peter Stone (Computer Science, College of Natural Sciences), left, and Virginia Tech's Sanmay Das explore the potential and challenges of using AI and machine learning to allocate scarce societal resources such as housing.

Highlights included an opening keynote from the London School of Economics' Robin Mansell examining the tensions between AI progress and democratic values, a cross-sector panel on autonomous vehicles and the future of transportation in smart cities featuring leaders from Waymo and the City of Austin and a closing keynote from Virginia Tech's Sanmay Das on the complexities of applying machine learning in high-stakes environments. Good Systems researchers showcased their latest work through presentations, live demos, posters and networking events.

Explore more

Celebrating a Defining Year for AI at UT Austin

UT Austin marked 2024 as the *Year of AI.* Good Systems helped drive this effort by leading cross-sector collaborative research, expanding its public programs and coordinating a national awareness campaign with partners throughout campus to share research and education resources with a broader audience. In partnership with city officials, Good Systems researchers developed AI tools to detect smoke and wildfires, improve emergency response times, make transportation safer and more accessible, increase transparency and stakeholder engagement and equip decision-makers with accurate data and predictive digital twin models. Good Systems-led initiatives also helped inform federal policy efforts around transparency and trust in AI and even helped advance ethics and safety in AI on the international stage.

Good Systems leaders and key partners share insights at AI LIVE, a three-day conference celebrating UT's Year of AI. From left: Chelsea McCullough (School of Information), Matt Lease (School of Information), Daniel Culotta (Chief Innovation Officer for the City of Austin), Junfeng Jiao (School of Architecture) and Luis Sentis (Aerospace Engineering and Engineering Mechanics, Cockrell School of Engineering).

The year culminated in AI LIVE, a campuswide celebration. The event featured a panel on Responsible AI Innovation with Good Systems leaders and community partners. The panel highlighted the importance of considering societal values in the development and deployment of AI technologies and discussed how values-driven innovation can advance equity and sustainability and improve quality of life for all.

Read more

UT Establishes Guidelines for AI in Education

In May, Good Systems and the Office of Academic Technology co-hosted a day-long campus event, Responsible AI for Education, bringing together voices from academia, industry and government to explore the future of AI in the classroom.

The event featured a series of panel discussions that included Good Systems researchers **Sharon Strover** and **Craig Watkins** (Journalism and Media, Moody College of Communication), **Sherri Greenberg** (LBJ School of Public Affairs), **Peter Stone** (Computer Science, College of Natural Sciences) and **Ken Fleischmann** (School of Information). The day concluded with the announcement of the **Responsible Adoption of AI Tools for Teaching and Learning framework**, a set of practical guidelines for using AI responsibly and ethically in educational settings. Created with input from faculty and staff across campus, the framework was open for public comment and is expected to be finalized by the start of the fall semester.

"Our collaboration with the Office of Academic Technology reflects a shared commitment to ensuring that AI tools are not only pedagogically effective but also ethical, transparent and aligned with human values."

—Sherri Greenberg, LBJ School of Public Affairs

From left: Peter Stone, Sherri Greenberg, Craig Watkins and Art Markman, senior vice provost for Academic Affairs, explore what responsible AI means in an educational context at the campus convening and announcement of UT's new Responsible Adoption of AI Tools for Teaching and Learning framework.

"Our collaboration with the Office of Academic Technology reflects a shared commitment to ensuring that AI tools are not only pedagogically effective but also ethical, transparent and aligned with human values," said Greenberg, a panel presenter and past Good Systems chair.

READ MORE

▶ **Watch the panel discussions**

Collaborations with Academic, Community and Industry Partners

Collaboration is at the heart of Good Systems' mission to design responsible AI systems for the benefit of society. It's essential to bring together perspectives across disciplines and sectors and from people from all walks of life and every corner of the globe. **Working with partners** from academia, industry, government and the nonprofit sector — including the City of Austin, MITRE, the Austin AI Alliance and Responsible AI UK — Good Systems researchers imagine and develop new AI systems and policies to advance equity and transparency, protect privacy and information integrity, equip and empower skilled workers, make our cities smarter and build robots that improve our lives.

Austin AI Alliance

Good Systems is a founding member of the Austin AI Alliance, a cross-sector nonprofit uniting over 120 organizations, including startups, tech companies, public agencies and academic institutions. This collaboration aims to shape a human-centered AI ecosystem in Central Texas, emphasizing education, policy, professional development and responsible innovation.

As a board member of the Austin AI Alliance, Good Systems' Sherri Greenberg, far left, moderated a panel discussion at the Alliance's State of AI in Austin event, exploring how to ethically and responsibly develop AI while keeping human intelligence at the forefront of innovation. Panelists included, from left, Michael Shepherd (Dell Technologies), Adam Klivans (Computer Science, College of Natural Sciences), Michael Wharton (KUNGFU.AI) and Lou Senko (Q2).

Good Systems' involvement focuses on translating ethical principles into actionable strategies. **Sherri Greenberg** (LBJ School of Public Affairs), former chair of Good Systems and current Alliance board member, plays a key role in guiding policy discussions and fostering community engagement. Notably, she moderated a panel at the Alliance's "State of AI in Austin" event, highlighting the importance of fairness and accountability in AI deployment.

"Fairness in development and deployment of AI — avoiding algorithmic discrimination, for example — is crucial," Greenberg said. "With generative AI, you're vacuuming up everything on the internet. If that includes biased or false information, it ends up embedded in the algorithms."

Read more

Peer Institution Collaborative

Good Systems took a leading role in launching a national "community of practice" dedicated to ethical and responsible AI research and education. The Alliance for Interdisciplinary AI (AIAI) brings together scholars from 10 leading institutions with interdisciplinary AI & Society programs — including Johns Hopkins, NC State, Rutgers, Penn State, Carnegie Mellon and others — to share best practices, develop shared resources for teaching and research in AI ethics and work together to foreground ethics in the national discourse around AI development, deployment and policy.

Spearheaded by the Good Systems Executive Team, this initiative reflects a shared commitment to interdisciplinary and multi-university collaboration in AI ethics leadership. **Chandra Bhat** (Maseeh Department of Civil, Architectural and Environmental Engineering, Cockrell School of Engineering), a founding member of Good Systems and part of its Executive Team, has been instrumental in shaping the collaborative's mission, emphasizing the importance of fostering critical thinking in students and encouraging them to question AI's capabilities and limitations. "We are in an age which has been compared to the Industrial Revolution or the Renaissance. It is fascinating and sometimes scary," Bhat said.

READ MORE

Leading a collaborative effort to create a community of practice with top universities, Good Systems' Stacey Ingram Kaleh, Sharon Strover (School of Journalism and Media, Moody College of Communication), Sam Baker (English, College of Liberal Arts), Sherri Greenberg (LBJ School of Public Affairs), Luis Sentis (Aerospace Engineering and Engineering Mechanics) and Chandra Bhat (Maseeh Department of Civil, Architectural and Environmental Engineering) were among those who convened in Washington, DC.

Pictured, from left: Kaleh, Shyam Sundar (Penn State University), Shaun Respess (NC State University), Qing Shen (University of Washington), Strover, Tony Dahbura (Johns Hopkins University), Baker, Ewa Ziarek (University of Buffalo), Mickey McGlasson (Carnegie Mellon University), Greenberg, Matthew Stone (Rutgers University), Lauren Goodlad (Rutgers University), Bhat, Sentis, Ivanka Pjesivac (University of Georgia), Lauren Tilton (University of Richmond), Debra Mathews (Johns Hopkins University), Jim Bellingham (Johns Hopkins University). Not pictured: Veljko Dubljevic (NC State University).

U.S. AI Safety Institute

To support its commitment to build safe, trustworthy AI systems, Good Systems joined the U.S. AI Safety Institute Consortium (now the Center for AI Standards and Innovation), a national effort led by the National Institute of Standards and Technology (NIST). As part of this collaboration, Good Systems partnered with The Cantellus Group and Good Tech Advisory LLC to host a two-day event focused on trust and transparency in AI.

Held in September on the UT Austin campus, the event convened leading voices from academia, government, industry and the nonprofit sector to explore a central question: How do we foster trust in AI amid disinformation, information overload and rapidly evolving synthetic media? The event featured a public roundtable with Councilwoman (now Mayor Pro Tem) Vanessa Fuentes (Austin City Council), Kay Firth-Butterfield (Good Tech Advisory), Steve Kramer (KUNGFU.AI) and Good Systems leader **Ken Fleischmann** (School of Information), followed by a private workshop with local C-suite leaders. Together, participants developed a set of guidelines for advancing trust and transparency in AI across sectors. A white paper synthesizing the workshop findings is in development.

Councilwoman (now Mayor Pro Tem) Vanessa Fuentes, left, and KUNGFU.AI's Steve Kramer participate in a public roundtable at Trusting AI: How Transparency Accelerates Human-AI Partnerships, a two-day event featuring a public roundtable and private workshop on UT campus that brought together leaders from academia, government, industry and the nonprofit sector to develop guidelines for advancing trust and transparency in AI.

Good Systems joined the U.S. AI Safety Institute Consortium, a national effort led by the National Institute of Standards and Technology (NIST), to foster trust and transparency in AI through multi-sector collaboration.

In January 2025, **Sam Baker** (Department of English, College of Liberal Arts) represented Good Systems at the Center for AI Standards and Innovation annual meeting. Baker shared Good Systems' interdisciplinary approach with AI experts from across the country, contributed to national conversations on building human-driven, stakeholder-centered AI systems and returned with insights from the workshop.

▶ **Watch the workshop video**

International Symposium on Trustworthy Autonomous Systems

Good Systems partnered with the UKRI Trustworthy Autonomous Systems Hub to co-host the second annual International Symposium on Trustworthy Autonomous Systems at UT Austin in September. The event, led by general chair **Justin Hart** (Computer Science, College of Natural Sciences), brought together more than 70 researchers from around the world to examine pressing questions at the intersection of ethics, AI and automation.

Through a series of workshops, talks and poster presentations, participants explored topics related to the use of autonomous systems in a range of settings

including transportation, healthcare, policing, defense, surveillance and sustainability. The symposium fostered interdisciplinary dialogue across fields such as computer science, philosophy, engineering, law, communications and social sciences. Good Systems researcher **Maria De-Arteaga** (Information, Risk and Operation Management, McCombs School of Business) delivered an invited keynote talk on limiting the role of machine learning in high-stakes contexts, and Good Systems Executive Team member **Sharon Strover** (School of Journalism and Media, Moody College of Communication) presented research on the impact of surveillance policies on public trust.

Researchers engage in interdisciplinary dialogue during the second annual symposium co-hosted by Good Systems and the UKRI Trustworthy Autonomous Systems Hub, which assembles teams from several UK universities to enable the development of socially beneficial autonomous systems.

Good Systems Headlines

External Features and Mentions

9/25/2024 **The Watson Wire: Innovation & AI in City Government** Watson Wire

10/9/2024 **What Could AI Do for Scientific Research?** GovTech

10/18/2024 **AI in Austin: How Central Texas officials are using it, and how they want to regulate it** CBS Austin

11/18/2024 **Parade of Robots takes over UT Austin campus to mark end of 'Year of AI'** Austin American-Statesman

12/31/2024 **As UT ends its Year of AI, experts create custom AI tutor to enhance learning for students** Austin American-Statesman

2/12/2025 **Austin AI Alliance sets out to measure progress of local efforts** Austin Business Journal

3/4/2025 **TribCast: The looming fight over regulating AI in Texas** Texas Tribune

3/12/2025 **Navigating the AI Landscape: Universities Unite for Ethical Innovation** Johns Hopkins University Institute for Assured Autonomy

UT Austin News Coverage

9/25/2024 **In the age of misinformation, seeing is not believing** The Daily Texan

10/16/2024 **As the presidential campaign goes viral, so, too, does misinformation** Moody College of Communication

11/22/2024 **UT Ranks 14th Globally for Interdisciplinary Science** UT News

12/24/2024 **UT Saw Huge Research Gains, Launched Education Initiatives During Its 'Year of AI'** UT News

5/20/2025 **UT Unveils Proposed Guidelines for Responsible Use of AI in Teaching and Learning** UT News

6/9/2025 **University releases proposed guidelines for artificial intelligence in teaching, learning, open to community feedback** Daily Texan

UT OVPR Communications

9/19/2024 **Three Researchers to Focus on Cross-Cutting Themes in Ethical AI**

9/19/2024 **Changing of the Good Systems Chairs**

9/25/2024 **Ethical Innovation: Good Systems Researchers and Partners on Shaping AI for Social Good**

11/1/2024 **Elections in the Age of AI: Good Systems Researchers Analyze AI's Role in Modern Political Campaigns**

12/12/2024 **Behind the Lens: Balancing Ethics and Innovation in Smart City Surveillance**

4/30/2025 **UT, Universities Join Forces on AI Ethics**

8/25/2025 **Unyielding Bias: How researchers are using video and VR to uncover how appearance affects pedestrian safety — and the risks of coding those patterns into AI**

Creating Connections

Good Systems continued to grow its dynamic network of interdisciplinary researchers and collaborative partners, both on campus and beyond. Faculty from nearly every college and school on campus contribute to the Good Systems network, helping to foster a culture of cross-disciplinary research collaboration.

Beyond the Forty Acres, Good Systems strengthened key relationships and deepened its involvement with local, national and international partners. These included the City of Austin, Austin AI Alliance, the Alliance for Interdisciplinary AI, MITRE and UKRI Trustworthy Autonomous Systems Hub (now Responsible AI UK).

The year also saw the formation of new connections. Good Systems joined the NIST U.S. AI Safety Institute Consortium and hosted a two-day convening with Cantellus Group and Good Tech Advisory. On campus, Good Systems joined forces with the Office of Academic Technology to bring university stakeholders together to explore responsible adoption and use of AI in teaching and learning and partnered with the Machine Learning Lab and Texas Robotics to form TexasAI — a central hub connecting the public to AI research, courses, programs and expertise at UT Austin. Good Systems also explored partnerships with industry leaders, including AMD's national Responsible AI team, Microsoft Research Dell's Community Impact team and IBM's global AI Ethics & Governance team.

As a trusted resource for policymakers, Good Systems continued to contribute to AI policy efforts at the city and state levels.

Explore these growing connections through our **interactive network map**. Search by name, college/school/unit or project, or click any node on the map and pause to see its connections appear. Zoom in for a closer look or select an individual to see which projects they've helped shape.

Funded Grants

To date, Good Systems researchers have secured more than $21 million in external funding through grants, gifts and awards that directly support the initiative's mission. External awards from the most recent fiscal year are listed below.

National Institutes of Health (National Library of Medicine)
Building Safety Guards into LLMs for Trustworthy Automatic Simplification of Medical Documents
$1,374,200
Jessy Li, College of Liberal Arts

Open Philanthropy
UT AI Alignment Initiative
$885,000
Brad Knox, College of Natural Sciences and Christian Tarsney, College of Liberal Arts

ExxonMobil
Domain-Targeted Language Models via Targeted Interventions
$110,000
Greg Durrett, College of Natural Sciences

Microsoft Research
Centering Ethics in the AI Curriculum: Scaling Up AI Ethics Education Nationwide
$80,000
Kenneth Fleischmann, School of Information

Knight Foundation
Data Enclaves for the Sharing of Sensitive Data
$68,000
Josephine Lukito, Moody College of Communication

City of Austin
AI Governance and Tools Training Curriculum Development for the City of Austin
$60,000
Sharon Strover, Moody College of Communication

Effective Altruism
Technical AI Research
$50,000
Brad Knox, College of Natural Sciences

Good Systems Partners

Good Systems collaborates with individuals and groups from various departments, organizations and institutions.

Internal

College of Natural Sciences	Machine Learning Lab
LBJ School of Public Affairs	Lab for Immersive Media
School of Architecture	Office of Academic Technology
Texas Robotics	UT Government Relations
IC2 Institute	Texas Development
Technology & Information Policy Institute	UT Austin Office of the President

External

AMD	MITRE
Austin AI Alliance	NC State University
Austin Community College	NIST U.S. AI Safety Institute Consortium
Austin Forum on Technology & Society	Penn State Center for Socially Responsible AI
Carnegie Mellon University	
City of Austin	Responsible AI UK
City of San Antonio	Rutgers University Critical AI
Chequeado	Smart Cities Connect
Full Fact	Texas Department of Information Resources (DIR)
Good Tech Advisory LLC	
Harvard University Berkman-Klein Center	Texas Innovation & Technology Caucus
Huston-Tillotson University	The Cantellus Group
Johns Hopkins University Institute for Assured Autonomy	UKRI Trustworthy Autonomous Systems Hub
MEASURE	University of Georgia
Meedan	University of Richmond
	University of Washington

Publications

Good Systems researchers advanced knowledge through successful multidisciplinary collaborations, leading to a wide range of peer-reviewed publications and conference presentations. The list below highlights the breadth of publications from the past fiscal year, most of which are available online.

*Barnard, P., Bautista, J. R., Dakhama, A., Farahi, A., Laos, K., Liu, A., & Menéndez, H. D. (2025). Responsible MLOps Design Methodology for an Auditing System for AI-Based Clinical Decision Support Systems. In H. D. Menéndez, G. Bello-Orgaz, P. Barnard, J. R. Bautista, A. Farahi, S. Dash, D. Han, S. Fortz, & V. Rodriguez-Fernandez (Eds.), *Testing Software and Systems* (Vol. 15383, pp. 217–236). Springer Nature Switzerland.

Biswas, J., Fussell, D., Stone, P., Patterson, K., Procko, K., Sabatini, L., & Xu, Z. (2025). The Essentials of AI for Life and Society: An AI Literacy Course for the University Community. *Proceedings of the AAAI Conference on Artificial Intelligence, 39*(28), 28973–28978.

Chandra, R., Zinage, V., Bakolas, E., Stone, P., & Biswas, J. (2025). Deadlock-free, safe, and decentralized multi-robot navigation in social mini-games via discrete-time control barrier functions. *Autonomous Robots, 49*(2).

Collier, C., Fleischmann, K. R., Greenberg, S. R., Lassiter, T., Kim, J. B., Salazar, C., New, R., & Longoria, R. G. (2024). Works for Me: Personalizing Skilled Trade Worker Training via Smart Hand Tools. *Proceedings of the Association for Information Science and Technology, 61*(1), 92–102.

Dao, T.-Q., Schneiders, E., Williams, J., Bautista, J. R., Seabrooke, T., Vigneswaran, G., Kolpekwar, R., Vashistha, R., & Farahi, A. (n.d.). *TAME Pain: Trustworthy AssessMEnt of Pain from Speech and Audio for the Empowerment of Patients* (Version 1.0.0) [Dataset]. PhysioNet.

Dao, T.-Q., Schneiders, E., Williams, J., Bautista, J. R., Seabrooke, T., Vigneswaran, G., Kolpekwar, R., Vashistha, R., & Farahi, A. (2025). TAME Pain data release: Using audio signals to characterize pain. *Scientific Data, 12*(1), 595.

Dass, S., Hu, J., Abbatematteo, B., Stone, P., & Martín-Martín, R. (2024, November). Learning to Look: Seeking Information for Decision Making via Policy Factorization. *Conference on Robot Learning*. Conference on Robot Learning (CORL 2024), Munich, Germany.

De-Arteaga, M., Jeanselme, V., Dubrawski, A., & Chouldechova, A. (2025). Leveraging Expert Consistency to Improve Algorithmic Decision Support. *Management Science*, mnsc.2022.01576.

Farahi, A., & Jiao, J. (2024). Analyzing racial disparities in the United States homeownership: A socio-demographic study using machine learning. *Cities, 152,* 105181.

*Francis, A., Pérez-D'Arpino, C., Li, C., Xia, F., Alahi, A., Alami, R., Bera, A., Biswas, A., Biswas, J., Chandra, R., Chiang, H.-T. L., Everett, M., Ha, S., Hart, J., How, J. P., Karnan, H., Lee, T.-W. E., Manso, L. J., Mirsky, R., … Martín-Martín, R. (2025). Principles and Guidelines for Evaluating Social Robot Navigation Algorithms. *ACM Transactions on Human-Robot Interaction, 14*(2), 1–65.

Govindarajan, V. S., Zang, M., Mahowald, K., Beaver, D., & Li, J. J. (2024). Do *they* mean 'us'? Interpreting Referring Expression variation under Intergroup Bias. *Findings of the Association for Computational Linguistics: EMNLP 2024*, 9772–9785.

Gupta, R., Shin, H., Norman, E., Stephens, K. K., Lu, N., & Sentis, L. (2024). Human Stress Response and Perceived Safety during Encounters with Quadruped Robots. *2024 33rd IEEE International Conference on Robot and Human Interactive Communication (ROMAN)*, 783–790.

Gupta, S., Kovatchev, V., Das, A., De-Arteaga, M., & Lease, M. (2025). Finding Pareto trade-offs in fair and accurate detection of toxic speech. *Information Research an International Electronic Journal, 30*(iConf), 123–141.

Hauser, E., Chan, Y.-C., Modak, S., Biswas, J., & Hart, J. (2024). Vid2Real HRI: Align video-based HRI study designs with real-world settings. *2024 33rd IEEE International Conference on Robot and Human Interactive Communication (ROMAN)*, 542–548.

Hejna, J., Rafailov, R., Sikchi, H., Finn, C., Niekum, S., Knox, W. B., & Sadigh, D. (2024). Contrastive Preference Learning: Learning from Human Feedback without RL. *The Thirteenth International Conference on Learning Representations.*

*Jiao, J., Seong, K., Sammer, M., Lewis, R. H., Reese, A., Olvera, N., Gronseth, S. L., Anderson-Fletcher, E., & Kakadiaris, I. (2025). Developing a healthy food access index (HFAI): Web-based mapping and future directions for AI integrations. *Cities, 161*, 105908.

Kambhatla, G., Lease, M., & Rajadesingan, A. (2024). Promoting Constructive Deliberation: Reframing for Receptiveness. *Findings of the Association for Computational Linguistics: EMNLP 2024,* 5110–5132.

*Kammer-Kerwick, M., Pogue, G., Purohit, I., Spandikow, E., De Jesus, R., Lifton, J., & Watkins, S. C. (2025). Integrating Health AI into Safety-Net and Rural Health Care Settings: Research, Insights, and Recommendations [Report Commissioned by the Episcopal Health Foundation].

Kolluri, A., Murthy, D., & Vinton, K. (2025). Quantifying the spread of racist content on fringe social media: A case study of Parler. *Big Data & Society, 12*(2), 20539517251321752.

Labiosa, A., Wang, Z., Agarwal, S., Cong, W., Hemkumar, G., Harish, A. N., Hong, B., Kelle, J., Li, C., Li, Y., Shao, Z., Stone, P., & Hanna, J. P. (2025, May). Reinforcement Learning Within the Classical Robotics Stack: A Case Study in Robot Soccer. International Conference on Robotics and Automation (ICRA).

*Li, H., Shin, H., Sentis, L., Siu, K.-C., Millán, J. D. R., & Lu, N. (2024). Combining VR with electroencephalography as a frontier of brain-computer interfaces. *Device, 2*(6), 100425.

Liu, H., Das, A., Boltz, A., Zhou, D., Pinaroc, D., Lease, M., & Lee, M. K. (2024). Human-centered NLP Fact-checking: Co-Designing with Fact-checkers using Matchmaking for AI. *Proceedings of the ACM on Human-Computer Interaction, 8*(CSCW2), 1–44.

Liu, H., Gwizdka, J., & Lease, M. (2025). Exploring Multidimensional Checkworthiness: Designing AI-assisted Claim Prioritization for Human Fact-checkers. *Proceedings of the ACM on Human-Computer Interaction (CSCW).*

Lukito, J. (2024a). Global Misinformation & Disinformation Special Issue Introduction. *International Journal of Public Opinion Research, 36*(3).

*Lukito, J. (2024b). Understanding a Case of State-sponsored Digital Disinformation. In *In Porter, L. (Ed.) The Disinformers: Social Media, Disinformation, and Elections.* Louisiana State University Press.

Lukito, J., Macdonald, M., Chen, B., Brown, M. A., Prochaska, S., Yang, Y., Greenfield, J., Suk, J., Zhong, W., Dahlke, R., & Borah, P. (2025). Candidates Be Posting: Multi-Platform Strategies and Partisan Preferences in the 2022 U.S. Midterm Elections. *Social Media + Society, 11*(2), 20563051251337541.

Mai, G., Xie, Y., Jia, X., Lao, N., Rao, J., Zhu, Q., Liu, Z., Chiang, Y.-Y., & Jiao, J. (2025). Towards the next generation of Geospatial Artificial Intelligence. *International Journal of Applied Earth Observation and Geoinformation, 136*, 104368.

Mao, L., Warnell, G., Stone, P., & Biswas, J. (2025). pacer: Preference-Conditioned All-Terrain Costmap Generation. *IEEE Robotics and Automation Letters, 10*(5), 4572–4579.

McCullough, C. E., Fleischmann, K. R., Greenberg, S. R., & Lassiter, T. B. (2025). Workers Who Care: AI-Enabled Smart Hand Tools in the Skilled Trades. *Proceedings of the ACM on Human-Computer Interaction, 9*(2), 1–20.

Muslimani, C., Chandramouli, S., Booth, S., Knox, B. W., & Taylor, M. E. (2024). Analyzing Reward Functions via Trajectory Alignment. *Proceedings of NeurIPS 2024 Workshop on Behavioral Machine Learning.* NeurIPS 2024 Workshop on Behavioral Machine Learning.

New, R., Salazar, C. D., Bendaña, J., Sripada V. S., S., Chinchali, S., Fleischmann, K. R., & Longoria, R. G. (2024). Design, Development, and Testing of a Smart Hand Tool: Achieving Work Task Recognition Using Synthetic Data and Edge Intelligence. *Volume 2A: 44th Computers and Information in Engineering Conference (CIE)*, V02AT02A026.

Norman, E., Gupta, R., Stephens, K., & Sentis, L. (2024, August). Utilizing Organizational Communication Theory for Community Embedded Robotics. *Proceedings of ALTRUIST 2024, BAILAR 2024, SCRITA 2024, and WARN 2024.* 33rd IEEE International Conference on Robot and Human Interactive Communication (RoMAN) - Workshop on Trust, Acceptance and Social Cues in Human-Robot Interaction (SCRITAS), Pasadena, CA.

Rafailov, R., Chittepu, Y., Park, R., Sikchi, H., Hejna, J., Knox, B., Finn, C., & Niekum, S. (2024). Scaling Laws for Reward Model Overoptimization in Direct Alignment Algorithms. *Proceedings of NeurIPS 2024 Workshop on Behavioral Machine Learning.* NeurIPS 2024 Workshop on Behavioral Machine Learning.

Romanow, J. (2024). Kipling's Administrative Imaginary: Imperial Ignorance as Information Management. *Victorian Studies, 66*(4), 585–608.

Ryan, G., Emily, N., Xingru, Z., Maria, E., Keri K., S., & Luis, S. (2025). Community Embedded Robotics: A Dataset to Study Perceived Social Intelligence and Safety During Unexpected Encounters with Quadrupedal Robots (Version 1.0) [Dataset]. Texas Data Repository.

Salazar, C. D., & Longoria, R. G. (2024). Comparative Analysis of Real-Time and Simulated Monitoring Techniques for MIG Welding. *Volume 2B: 44th Computers and Information in Engineering Conference (CIE),* V02BT02A036.

*Sammer, M., Seong, K., Olvera, N., Gronseth, S. L., Anderson-Fletcher, E., Jiao, J., Reese, A., & Kakadiaris, I. A. (2024). AI-FEED: Prototyping an AI-Powered Platform for the Food Charity Ecosystem. *International Journal of Computational Intelligence Systems, 17*(1), 259.

Schoeffer, J., De-Arteaga, M., & Elmer, J. (2025, April 5). Perils of Label Indeterminacy: A Case Study on Prediction of Neurological Recovery After Cardiac Arrest. *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency (FAccT).*

Seo, M., Cho, Y., Sung, Y., Stone, P., Zhu, Y., & Kim, B. (2025, May). PRESTO: Fast Motion Planning Using Diffusion Models Based on Key-Configuration Environment Representation. IEEE International Conference on Robotics and Automation (ICRA).

Seong, K., Jiao, J., Mandalapu, A., & Niyogi, D. (2024). Spatio-temporal patterns of heat index and heat-related Emergency Medical Services (EMS). *Sustainable Cities and Society, 111*, 105562.

Sriram, A., Xu, F., Choi, E., & Durrett, G. (2024, November 12). Contrastive Learning to Improve Retrieval for Real-world Fact Checking. The 2024 Conference on Empirical Methods in Natural Language Processing, Miami, Florida.

Stark, C., Chun, B., Charleston, C., Ravi, V., Pabon, L., Sunkari, S., Mohan, T., Stone, P., & Hart, J. (2024). Dobby: A Conversational Service Robot Driven by GPT-4. *2024 33rd IEEE International Conference on Robot and Human Interactive Communication (ROMAN)*, 1362–1369.

Strover, S., Lalwani, S., & El-Masri, A. (2024). Uneven Eyes: The Impact of Inconsistent Local Surveillance Policies on Public Trust. *Proceedings of the Second International Symposium on Trustworthy Autonomous Systems*, 1–5.

Su, Y., Li, J. J., & Lease, M. (2024). Wrapper Boxes for Faithful Attribution of Model Predictions to Training Data. *Proceedings of the 7th BlackboxNLP Workshop: Analyzing and Interpreting Neural Networks for NLP*, 551–576.

Sui, P., Rodriguez, J. D., Laban, P., Murphy, D., Dexter, J. P., So, R. J., Baker, S., & Chaudhuri, P. (2025, July). KRISTEVA: Close Reading as a Novel Task for Benchmarking Interpretive Reasoning. Association for Computational Linguistics, Vienna, AU.

Tang, C., Abbatematteo, B., Hu, J., Chandra, R., Martín-Martín, R., & Stone, P. (2025). Deep Reinforcement Learning for Robotics: A Survey of Real-World Successes. *Annual Review of Control, Robotics, and Autonomous Systems, 8*(1), 153–188.

Tang, L., Laban, P., & Durrett, G. (2024). MiniCheck: Efficient Fact-Checking of LLMs on Grounding Documents. *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing,* 8818–8847.

Verzic, N., Chadaga, A., & Hart, J. (2024). Recovering Missed Detections in an Elevator Button Segmentation Task. *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 13355–13362.

Williams, J., Schneiders, E., Card, H., Seabrooke, T., Pakenham-Walsh, B., Azim, T., Valls-Reed, L., Vigneswaran, G., Bautista, J. R., Chandra, R., & Farahi, A. (2024). Predicting Acute Pain Levels Implicitly from Vocal Features. *Interspeech 2024*, 1460–1464.

Xu, Y., Jiao, J., & Wang, H. (2025, January 5). An Urban Digital Twin Framework for Sustainable Transportation and Smart Cities: A Case Study of Austin, TX. 104th Transportation Research Board (TRB) *Annual Meeting. 104th Transportation Research Board (TRB) Annual Meeting,* Washington, D.C.

Yang, H., Lee, J.-W., Klepac, S., Santos Cruz, A. U., Subgranon, A., & Jiao, J. (2024). Reconstructing missing data of damaged buildings from post-hurricane reconnaissance data using XGBoost. *Frontiers in Built Environment, 10*, 1444001.

Zhang, A., & Lee, M. K. (2025). Knowledge Workers' Perspectives on AI Training for Responsible AI Use. *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 1–18.

Zhang, A., Liao, M., Elizaveta, Kravchenko, Taylor, M., Haddad, A., Bhat, C., Watkins, S. C., & Lee, M. K. (2025). Data and Technology for Equitable Public Administration: Understanding City Government Employees' Challenges and Needs. *Proceedings of the ACM on Human Computer Interaction (HCI).*

Zhang, M. J. Q., Knox, W. B., & Choi, E. (2025, April). Modeling Future Conversation Turns to Teach LLMs to Ask Clarifying Questions. The Thirteenth International Conference on Learning Representations, Singapore.

* Indicated publications do not formally acknowledge Grand Challenge support but are verified Good Systems research project outputs.

Good Systems Team

Executive Team

Good Systems' Executive Team represents the College of Liberal Arts, the College of Natural Sciences, the Moody College of Communication, the Cockrell School of Engineering, the LBJ School of Public Affairs, the School of Architecture and the School of Information.

Luis Sentis
Chair (2024-25)
Aerospace Engineering and Engineering Mechanics
Cockrell School of Engineering

Samuel Baker
Past Chair (2021-22) and Founding Member
English
College of Liberal Arts

Chandra Bhat
Founding Member
Maseeh Department of Civil, Architectural and Environmental Engineering
Cockrell School of Engineering

Kenneth R. Fleischmann
Founding Chair
School of Information

Sherri Greenberg
Past Chair (2023-24)
LBJ School of Public Affairs

Junfeng Jiao
Past Chair (2020-21) and Founding Member
School of Architecture

Matthew Lease
Founding Member
School of Information

Peter Stone
Founding Member
Computer Science
College of Natural Sciences

Sharon Strover
Past Chair (2022-23) and Founding Member
School of Journalism and Media
Moody College of Communication

Lea Sabatini
Program Director
Office of the Vice President for Research, Scholarship and Creative Endeavors

Stacey Ingram Kaleh
Associate Director, Partnerships and Programs
Office of the Vice President for Research, Scholarship and Creative Endeavors

Core Research Projects

Designing AI to Advance Racial Equity

S. Craig Watkins
School of Journalism and Media
Moody College of Communication
Project Lead

Chandra Bhat
Maseeh Department of Civil, Architectural and Environmental Engineering
Cockrell School of Engineering
Project Co-Lead

DeLawnia Comer-HaGans
Office of Health Equity
Dell Medical School

Kaya de Barbaro
Psychology
College of Liberal Arts

Min Kyung Lee
School of Information
Project Co-Lead

Matt Kammer-Kerwick
IC^2 Institute

Alex Karner
School of Architecture

Lisa Macias
Maseeh Department of Civil, Architectural and Environmental Engineering
Cockrell School of Engineering

Kenneth Perrine
Center for Transportation Research
Cockrell School of Engineering

Ruben Rathnasingham
Medical Education
Dell Medical School

Sharon Ricks
Office of Health Equity
Dell Medical School

Maytal Saar-Tsechansky
Information, Risk and Operations Management
McCombs School of Business

Being Watched: Embedding Ethics in Public Cameras

Sharon Strover
School of Journalism and Media
Moody College of Communication
Project Lead

Atlas Wang
Chandra Family Department of Electrical
and Computer Engineering
Cockrell School of Engineering
Project Co-Lead

Terrence Chapman
Government
College of Liberal Arts

Joydeep Ghosh
Chandra Family Department of Electrical
and Computer Engineering
Cockrell School of Engineering

Junyuan Hong
Chandra Family Department of Electrical
and Computer Engineering
Cockrell School of Engineering

Brad Limov
School of Journalism and Media
Moody College of Communication

Amy Sanders
School of Journalism and Media
Moody College of Communication

Anita Varma
School of Journalism and Media
Moody College of Communication

Making Smart Tools Work for Everyone

Kenneth R. Fleischmann
School of Information
Project Lead

Sherri Greenberg
LBJ School of Public Affairs
Project Co-Lead

Raul Longoria
Walker Department of Mechanical
Engineering
Cockrell School of Engineering
Project Co-Lead

Sandeep Chinchali
Chandra Family Department of Electrical
and Computer Engineering
Cockrell School of Engineering

Ashish Deshpande
Walker Department of Mechanical
Engineering
Cockrell School of Engineering

Zhenghui Sha
Walker Department of Mechanical
Engineering
Cockrell School of Engineering

Designing Responsible AI Technologies to Protect Information Integrity

Matthew Lease
School of Information
Project Lead

Dhiraj Murthy
School of Journalism and Media
Moody College of Communication
Project Co-Lead

Maria De-Arteaga
Information, Risk and Operations
Management
McCombs School of Business

Greg Durrett
Computer Science
College of Natural Sciences

Jacek Gwizdka
School of Information

Min Kyung Lee
School of Information

Jessy Li
Linguistics
College of Liberal Arts

Josephine "Jo" Lukito
School of Journalism and Media
Moody College of Communication

Ashwin Rajadesingan
Communication Studies
Moody College of Communication

Maytal Saar-Tschansky
Information, Risk and Operations
Management
McCombs School of Business

A Good System for Smart Cities

Junfeng Jiao
School of Architecture
Project Lead

Arya Farahi
Statistics and Data Sciences
College of Natural Sciences
Project Co-Lead

Dev Niyogi
Earth and Planetary Sciences
Jackson School of Geosciences
Maseeh Department of Civil,
Architectural and Environmental
Engineering
Cockrell School of Engineering
Project Co-Lead

Yuhao Kang
Geography and the Environment
College of Liberal Arts

Jun-Whan Lee
Maseeh Department of Civil,
Architectural and Environmental
Engineering
Cockrell School of Engineering

Gengchen Mai
Geography and the Environment
College of Liberal Arts

Kijin Seong
School of Architecture

Zhenghui Sha
Walker Department of Mechanical
Engineering
Cockrell School of Engineering

Atlas Wang
Chandra Family Department of Electrical
and Computer Engineering
Cockrell School of Engineering

Junmin Wang
Walker Department of Mechanical
Engineering
Cockrell School of Engineering

Yiming Xu
School of Architecture

Living and Working with Robots

Samuel Baker
English
College of Liberal Arts

Joydeep Biswas
Computer Science
College of Natural Sciences

Samanta Varela Castro
Communication Studies
Moody College of Communication

Boh Chun
School of Information

Aaron Choate
University of Texas Libraries

Maria Esteva
Texas Advanced Computing Center

Justin Hart
Computer Science
College of Natural Sciences

Elliott Hauser
School of Information

Junfeng Jiao
School of Architecture

Adam Klivans
Computer Science
College of Natural Sciences

Brad Knox
Computer Science
College of Natural Sciences

Roberto Martín-Martín
Computer Science
College of Natural Sciences

Katie Pierce Meyer
University of Texas Libraries

Desmond Ong
Psychology
College of Liberal Arts

Jacob Romanow
English
College of Liberal Arts

Luis Sentis
Aerospace Engineering and Engineering
Mechanics
Cockrell School of Engineering

Samantha Shorey
Communication Studies
Moody College of Communication

Keri Stephens
Communication Studies
Moody College of Communication

Peter Stone
Computer Science
College of Natural Sciences

Yuke Zhu
Computer Science
College of Natural Sciences

Cross-Cutting Themes

Alignment and Safety

Brad Knox
Computer Science
College of Natural Sciences

Knowledge, Generative AI and Power

Jared Jensen
School of Journalism and Media
Moody College of Communication

Work at the Human-AI Frontier

Lingyuan Li
School of Information