


GOOD SYSTEMS

FISCAL YEAR 2026 ANNUAL REPORT



TEXAS

The University of Texas at Austin
Office of the Vice President for Research

GOOD SYSTEMS

FISCAL YEAR 2026 ANNUAL REPORT

Table of Contents

Good Systems Highlights in Numbers3

Good Systems Research4

Core Projects5

Cross-Cutting Themes14

Seed Projects16

Creating Connections18

Sharing the Story of Good Systems.19

Research Engagement19

Collaborations with Academic, Community
and Industry Partners.23

Good Systems Partners26

Good Systems Headlines27

Funded Grants28

Publications29

Good Systems Team32



ADVANCING ETHICAL, TRUSTWORTHY INNOVATION IN AI IS OUR GRAND CHALLENGE.

Artificial intelligence now shapes how people work, move through their cities, find information and connect with one another, often in ways that go unnoticed until something goes wrong. Its benefits are real, but so are its risks, and the two are rarely easy to separate. Good Systems operates from the conviction that designing AI technologies well is not only a technical problem but a human one, inseparable from the values, livelihoods and rights of the people these systems touch.

Based at The University of Texas at Austin, Good Systems is a campuswide research grand challenge that brings technologists, humanists and social scientists together to define, evaluate and build ethical AI systems. Its teams work together with city and state agencies, industry, nonprofits and community members, keeping the people affected by AI at the center of how those systems are designed.

Underlying these efforts is a shared definition: a “good” system is a human-AI partnership that strengthens rather than erodes human agency, fosters trust and transparency, and supports the foundations of an open and participatory society.

As Good Systems looks toward its next chapter, **Kenneth Fleischmann** (School of Information), who chairs the Good Systems Executive Team, said the grand challenge has only reinforced the ongoing need for this work. “I’m most optimistic about — and most confident about — the degree to which we chose the right problem,” he said. “This is an important societal challenge.”

“I’m most optimistic about — and most confident about — the degree to which we chose the right problem. This is an important societal challenge.”

**KENNETH FLEISCHMANN,
SCHOOL OF INFORMATION**

GOOD SYSTEMS HIGHLIGHTS IN NUMBERS

Expanding Networks

66 active
researchers

33 UT departments
& disciplines

16 schools &
colleges

30 external
partners

11 events hosted
over the past year

Engaging Students

77 undergraduate &
graduate student
researchers



Scholarly Output & Publicity

32 scholarly works published
over the past year
(acknowledging
Good Systems support)

19 news articles over
the past year

Building Capacity

\$52.75 M

awarded in external
funding to date



GOOD SYSTEMS RESEARCH

Good Systems' research has matured from defining what ethical AI requires to building and field-testing the tools, methods and frameworks that put those ideas to work. In Fiscal Year 2026, the grand challenge continued research across its six core projects and two cross-cutting themes and initiated a new cohort of seed-funded teams that extend its work into areas such as maternal health, flood resilience and disability access. The projects pair technical advances with the policy and design questions that determine how AI technologies are actually used — frequently in partnership with the city agencies, workers and communities they affect.

The highlights that follow reflect that mix of technical depth and applied, real-world engagement.



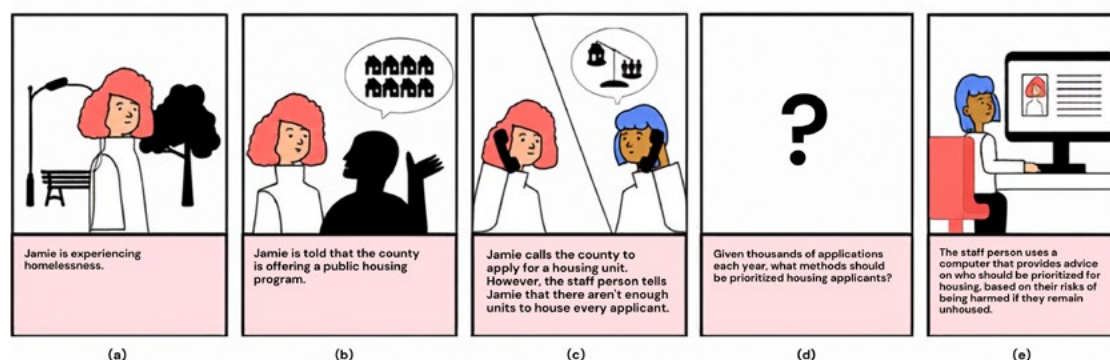
Core Projects

BEING WATCHED: EMBEDDING ETHICS IN PUBLIC CAMERAS

How should public agencies govern the AI technologies they deploy? That question has become more urgent as municipalities across the country have integrated powerful new tools — from automated license plate readers and drone cameras to large language model applications — into routine operations. This project continued to develop answers through technical innovation, applied research and direct partnership with local government.

Project leader **Sharon Strover** (School of Journalism and Media, Moody College of Communication) led a collaboration with the City of Austin to design and deliver AI literacy training for city employees. Supported by funding through the Interlocal Agreement between UT and the City of Austin, the work was animated by a core conviction: equipping public servants to think critically about AI is as important as deploying the tools themselves.

Working with postdoctoral researcher **Brad Limov** and doctoral student **Azza El-Masri** (both from the School of Journalism and Media, Moody College of Communication), Strover developed an ethics-first curriculum that asked participants to examine the harder questions — Who is accountable if something goes wrong? What data are being used? Who might be affected? — before reaching for a tool. The team surveyed approximately 1,500 city employees, conducted six interactive workshops with more than 100 participants across departments and produced a final report for the city along with a 10-part video series on responsible generative AI use.



Captured from the eighth video (“Societal Harms and Environmental Impacts”) in a 10-part series the team produced for the City of Austin. Project co-leader Brad Limov (bottom left) presents this unit.



The workshops engaged a workforce that was both curious and uneven in its familiarity with AI. “There were people who had never used AI, who had never even thought about ChatGPT, and then there were people who were using it daily,” said El-Masri, who helped lead the sessions. That range shaped the training itself: case studies were revised based on participant feedback as employees clarified how their work actually functioned in practice.

[READ MORE](#)

The Being Watched project also broadened its technical privacy research from camera systems into large language models, vision-language models and AI agents. [Atlas Wang](#) and [Jianing Zhu](#) (both from the Chandra Family Department of Electrical and Computer Engineering, Cockrell School of Engineering) continued work on privacy benchmarking, agent safety and machine unlearning, a set of methods that allows AI models to selectively forget learned concepts or sensitive data without degrading overall performance. The work is especially relevant for public-sector AI systems, where agencies need tools that can protect privacy, support accountability and respond to demands for data removal after systems are deployed.

As national debates over AI’s role in government intensified, Strover contributed directly to those conversations. Her essay in Tech Policy Press examined the limits of corporate ethical commitments in the context of the Anthropic-Pentagon dispute, arguing that durable accountability requires democratic structures, not just the principled stands of individual companies. The piece brought the project’s central concerns into a national policy debate, emphasizing that AI systems used in surveillance and public safety require public governance, not just technical safeguards or corporate restraint.

[READ MORE](#)

DESIGNING RESPONSIBLE AI TECHNOLOGIES TO PROTECT INFORMATION INTEGRITY

As AI systems take on greater roles in generating, summarizing and evaluating information, the question of whether those systems can be trusted has become central to research on the information environment. In FY26, the project deepened its focus on that question across three converging lines of inquiry: the reliability of AI outputs, the trustworthiness of large language models in high-stakes settings and the ways context and culture shape how people engage with AI-mediated information.

Those questions shape work by **Matthew Lease** (School of Information), who studied how users can know how much confidence to place in an AI-generated answer. Lease investigated quality estimation and uncertainty quantification for free-form generative AI outputs, developing methods that predict output quality and estimate ranges of confidence around those predictions. The work matters because many AI systems produce answers that sound authoritative even when their factual basis is uncertain; better quality estimates could help users decide when outputs require additional scrutiny.

Lease also presented two papers at the 2026 Association for Computing Machinery (ACM) Conference on Fairness, Accountability, and Transparency that examined trust and fairness in AI-assisted information work from different angles. In one study, led by former postdoctoral fellow **Sanjana Gautam** with doctoral students **Houjiang Liu** and **Yujin Choi** (all from the School of Information), participants described how they decide whether to trust AI tools in early-stage research. The study found that users often struggled to tell when AI-generated answers required closer review, especially when the system's wording sounded confident or the source of its information was difficult to trace.

Jessy Li (Linguistics, College of Liberal Arts) continued her research on the safety boundaries of large language models in medical contexts. Her team found that LLMs will provide potentially dangerous clinical answers when false information is present in the prompt — even when the model has access to accurate information — raising important questions about the conditions under which these tools can be safely used. The findings were presented at the Annual Meeting of the Association for Computational Linguistics.

The project also examined how AI-mediated information is shaped by social and cultural context. **Dhiraj Murthy** (School of Journalism and Media, Moody College of Communication) studied information integrity and AI trust in the Global South, finding that AI use is shaped by existing social practices, media habits and local contexts rather than following a single pattern of disruption. His [article](#) in *New Media & Society*, based on workplace interviews in India, has implications for how researchers and policymakers assess the risks and governance needs of AI systems in diverse societies.

In professional fact-checking, **Min Kyung Lee** (School of Information) studied how media experts evaluate AI-generated fact-checking reports.

The work matters because many AI systems produce answers that sound authoritative even when their factual basis is uncertain; better quality estimates could help users decide when outputs require additional scrutiny.

Her team found that experts often regarded AI-produced reports as accurate but less useful than those authored by humans, a distinction with implications for how AI tools are integrated into editorial and analytical workflows. The work was presented at the 2026 ACM Computer-Human Interaction Conference on Human Factors in Computing Systems.

Finally, **Jo Lukito** (School of Journalism and Media, Moody College of Communication) examined cross-platform election communication and misinformation using large-scale datasets, including **Candidata**, a public archive and research dataset developed by the Algorithmic Transparency Institute that contains social media data, handles and posts for U.S. federal election candidates. She presented **findings** at the International Association for the Advancement of AI Conference on Web and Social Media. Lukito, a contributor to the project since its inception, departed UT for the University of Southern Denmark in January 2026.

LIVING AND WORKING WITH ROBOTS


In FY26, Living and Working with Robots helped shape the research agenda for a new \$30 million National Science Foundation Science and Technology Center focused on co-adaptive robotics and community engagement. Led by **Joydeep Biswas** (Computer Science, College of Natural Sciences), the center will build on Good Systems research by studying long-term robot deployments in ordinary settings, including shared campus spaces, and examining how robots and human users adapt to one another over time. That approach will generate more ecologically valid data on user experience, robot architectures and human-robot interaction than short-term lab studies can typically provide.

READ MORE

In other research, Good Systems' multiyear study on indoor human-robot encounters reached a scholarly milestone with publication in ACM Transactions on Human-Robot Interaction. Led by **Keri Stephens** and **Emily Norman** (both from the Department of Communication Studies, Moody College of Communication), **Maria Esteva** (Texas Advanced Computing Center), **Hyonyoung Shin** (Chandra Family Department of Electrical and Computer Engineering, Cockrell School of Engineering), and **Nanshu Lu, Luis Sentis** and **Ryan Gupta** (all from Aerospace Engineering and Engineering Mechanics, Cockrell School of Engineering), the transdisciplinary study examined perceived safety in human-robot interactions: how comfortable and secure people feel when mobile robots share the spaces where they live, work and gather. The resulting framework can help researchers design

robotic systems whose movement and behavior support both physical safety and a greater sense of comfort in shared environments. Their [article](#) was also highlighted in Science Robotics as an Editors' Choice, recognizing its contribution to an aspect of human-robot interaction that has received less attention, particularly the experiences of people who are near robots without directly interacting with them.

Meanwhile, technical research on social navigation, perception and robotic learning continued. Biswas, [Justin Hart](#), [Roberto Martín-Martín](#) and [Peter Stone](#) (all from the Department of Computer Science, College of Natural Sciences) published findings on semantic mapping, human motion prediction and reinforcement learning for robot control at the Conference on Robot Learning and the IEEE/Robotics Society of Japan International Conference on Intelligent Robots and Systems. Research published in The International Journal of Robotics Research introduced [LLM-GROP](#), an approach integrating large language models with task and motion planning so robots can translate spoken or written service requests into actions based on the objects and spatial relationships in front of them.




An illustration of the mobile manipulation domain, where a mobile manipulator is tasked with setting a dining table — from “LLM-GROP: Visually Grounded Robot Task and Motion Planning with Large Language Models,” Zhang et al.’s article in The International Journal of Robotics Research. The robot must arrange several tableware items, including a knife, a fork, a plate, a teacup and a strawberry. These objects are located on other tables, and the environment also includes randomly generated obstacles (e.g., chairs), which are not accounted for in the pre-built map.

SMART HAND TOOLS: BUILDING THE FUTURE OF WORK WITH ETHICAL AI

The core tenet of this project is that the future of skilled work lies not in replacing human expertise but in augmenting it. Focused on the middle-skill trades, the team develops smart hand tools through participatory design: studies with workers and apprentices inform the design of AI-enabled sensor attachments, and subsequent studies examine how those workers understand the tools' implications for safety, skill, surveillance and data ownership.

Kenneth Fleischmann (School of Information) and **Sherri Greenberg** (LBJ School of Public Affairs), along with graduate student co-authors **Chelsea McCullough**, **Jessica Needle** and **Haley Triem** (all from the School of Information), surveyed and interviewed unionized skilled trade workers and apprentices to understand their attitudes toward AI, workplace surveillance, safety and data ownership. Their article, **“AI Is Not Gonna Take Our Jobs,”** was published in IEEE Transactions on Technology and Society, and their findings ran counter to a common assumption about automation anxiety. Participants were generally not worried about AI replacing their jobs; rather, they identified safety as the most valuable use case for AI-enabled tools and expressed a strong desire to retain control over the data those tools generate.



Captured from “Designing AI Systems for Good: Collaborative Research with Transformative Impact,” a video Good Systems released in spring 2026 about the impact of the grand challenge and the continuing need for ethical AI research. In the video, Carlos Salazar demonstrates the welding simulator he designed and developed with the Smart Hand Tools project team.

On the engineering side, **Raul Longoria** and graduate students **Carlos Salazar** and **Rachel New** (all from the Walker Department of Mechanical Engineering, Cockrell School of Engineering) continued developing a modular sensor unit that can be attached to conventional power tools. The prototype is designed to collect data about how a tool is used, then process those data directly on the device rather than relying on remote analysis.

The team used the prototype to test whether smart hand tools can recognize tasks, assess skills and detect signs of unusual tool performance in real time. The **results** — presented at the American Society of Mechanical Engineers 2025 International Design Engineering Technical Conferences & Computers and Information in Engineering Conference — showed the feasibility of real-time, on-device analytics that could eventually give workers feedback during tool operation and generate post-work reports without removing the worker from the center of the task.

The project also looked toward its educational potential. Joined by **Zhenghui Sha** (Walker Department of Mechanical Engineering, Cockrell School of Engineering), the team is exploring how the project's participatory design datasets can support engineering education. Drawing on data from partners such as the City of Austin and Longhorn Racing, the team is developing an LLM-based platform with which students can interact with virtual stakeholder agents, helping them learn how to interpret stakeholder needs and elicit requirements during the conceptual design process.

DESIGNING AI FOR SOCIAL FAIRNESS

As AI continues to move into transportation systems, public agencies and clinical settings, fairness depends not only on how models are built but also on whose experiences, data and needs shape them. In FY26, Designing AI for Social Fairness focused on that problem within several domains, including pedestrian safety, public-sector decision-making and health AI. Across those areas, the project paired technical research with studies of the social and institutional conditions that determine whether AI systems can be used responsibly.

In transportation, the team continued its study of pedestrian safety and the factors that shape pedestrian crossing and driver yielding behaviors. After completing analysis of more than 1,000 hours of intersection footage, a team led by **Chandra Bhat** (Maseeh Department of Civil, Architectural and Environmental Engineering, Cockrell School of Engineering) used virtual reality to study crossing behaviors under controlled conditions, conducting hour-long sessions with more than 80 participants. The project examines

A decorative header image featuring a complex circuit board pattern with glowing blue and white lines on a dark background.

how situational factors, such as traffic volume, signal timing and the behavior of nearby pedestrians, interact with individual characteristics to shape crossing decisions. The team published **findings** from its video analysis in Accident Analysis & Prevention. The larger goal is to inform the development of fairer AI models, such as autonomous vehicle detection algorithms, so that they recognize a complete range of pedestrian behaviors and omit biases from driver-pedestrian interactions.

The project also expanded its research on AI applications in healthcare, led by **S. Craig Watkins** (School of Journalism and Media, Moody College of Communication). One study, developed with **Kaya de Barbaro** (Psychology, College of Liberal Arts), examined the social, technical and cultural barriers that limit participation by diverse mothers in research using smartphones, wearables and other sensor-based tools to understand maternal health and caregiving practices. Interviews with health professionals and mothers revealed that trust depends not only on what the technology can measure but on whether participants understand how data will be collected, who can access it and whether they can control or withdraw from the research.


A second line of health AI work moved toward clinical deployment. The team worked with the kidney transplant team at Dell Medical School to scope a pilot digital health platform designed to support remote monitoring of recovery after transplant. Through weekly meetings with clinicians, the researchers refined the platform around clinical workflow, patient needs, data architecture and dashboard design. The goal is to capture behavioral, psychosocial, medical and demographic data in ways that help clinicians better understand recovery trajectories while testing whether the system is feasible, usable and acceptable in a clinical setting.

In public-sector technology, **Min Kyung Lee** (School of Information) continued research on how government organizations use data and AI to promote equity. Her team's paper, **"Data and Technology for Equitable Public Administration,"** received a Best Paper Honorable Mention at the Association for Computing Machinery Conference on Computer-Supported Cooperative Work and Social Computing. Drawing on interviews with city government employees, the work identifies the data, tools and organizational support they need to pursue equity goals. In related work, the team developed Contextualizing Datasets, a tool that helps public-sector employees examine datasets for subjectivity, bias and gaps in representation before those data are used in decision-making or generative AI systems. The work was presented at the 2026 CHI Conference on Human Factors in Computing Systems.

A GOOD SYSTEM FOR SMART CITIES

As cities integrate AI, it is not just about what these systems can do but whom they are designed to serve. Using the City of Austin as a case study, this project designs and deploys AI tools, including digital twin models, to support fairer, better-informed decision-making when it comes to city planning and public services.

Building on the project's digital twin infrastructure, project leader **Junfeng Jiao**, postdoctoral fellow **Yiming Xu** and graduate student **Jihyung Park** (all from the School of Architecture) developed a web-based transportation simulation platform that combines SUMO, a widely used traffic-modeling system, with LLM agents. Users can describe a planning question or proposed intervention in everyday language, and the system translates that request into changes to the simulation, such as retiming traffic signals, modifying road networks or modeling special events. By making detailed traffic simulations easier for nonexperts to use, the platform could help planners, policymakers and researchers test mobility strategies and evaluate urban resilience. The team presented the work at the [NeurIPS UrbanAI Workshop](#).



An LLM-powered transportation digital twin developed by Yiming Xu, Jihyung Park and Junfeng Jiao. The web-based platform allows users to describe transportation scenarios and proposed changes in everyday language, which the system translates into changes in a simulation of Austin's transportation network.

The team built a web-based simulation of downtown Austin that uses a physics-constrained deep learning framework to estimate travel times from medical centers and fire stations to emergency locations. The results show how digital twin-based modeling can help identify practical ways to improve EMS response times. The team is also developing a performance dashboard with spatial response metrics to support local stakeholders, such as Austin-Travis County EMS, in evaluating alternative response strategies under various urban disruptions.

The project's digital-twin and urban AI research also extended into mobility analytics and hazard response. **Gengchen Mai** (Geography and the Environment, College of Liberal Arts) co-authored a study in the **International Journal of Geographical Information Science** introducing a knowledge-driven GeoAI framework for classifying urban e-scooter ridership, connecting geospatial AI with street-level infrastructure and mobility patterns. **Jun-Whan Lee** (Maseeh Department of Civil, Architectural and Environmental Engineering, Cockrell School of Engineering) co-authored a study in **Journal of Hydrology: Regional Studies** using machine learning-enhanced flood models to estimate high-resolution storm surge inundation in Southeast Texas. Together, the studies show how the project's AI-enabled urban modeling work can support both everyday mobility planning and preparedness for environmental hazards.

Cross-Cutting Themes

Building on the six core projects, Good Systems researchers deepened work through two cross-cutting themes that connect efforts across teams and support new lines of inquiry.

AI ALIGNMENT AND SAFETY

Brad Knox (Computer Science, College of Natural Sciences), who leads the Alignment and Safety cross-cutting theme, focused much of his research on the design and risks of AI companions. Collaborating with **Samuel Baker** (English, College of Liberal Arts), **Desmond Ong** (Psychology, College of Liberal Arts) and **Sean Williams** (School of Law), Knox developed a systematic framework for understanding the potential harms of AI companions. Their research identifies nearly 50 causal pathways through which AI companion systems can harm users, their relationships and society, from fostering unhealthy attachment to reducing engagement

with the physical world. Knox shared the framework in several invited talks, including the University of Toronto's Schwartz Reisman Institute for Technology and Society Seminar Series.

[READ MORE](#)

Knox's work attracted wide attention, including a [feature in IEEE Spectrum](#), where he described the urgency of building this conceptual vocabulary before the harms become entrenched. "The more we can quickly develop a sophisticated understanding of what AI companions are doing to their users, to their users' relationships and to society at large," Knox said, "the sooner we can apply that understanding to their design."

Knox also launched a new study with the team — a representative survey of roughly 1,000 adolescents on their use of AI for social-emotional purposes. He also co-leads the [AI+Human Objectives Initiative](#), which hosts a series of interdisciplinary talks on AI alignment, safety and social impact at UT Austin.

[READ MORE](#)

"The more we can quickly develop a sophisticated understanding of what AI companions are doing to their users, to their users' relationships and to society at large, the sooner we can apply that understanding to their design."

**BRAD KNOX,
COLLEGE OF
NATURAL SCIENCES**

KNOWLEDGE, GENERATIVE AI AND POWER

Leading the Knowledge, Generative AI and Power theme, **Jared Jensen** (Communication Studies, Moody College of Communication) continued his research into two related lines of inquiry: how generative AI is reshaping creativity and the creative industries, and how AI ethics are enacted within interdisciplinary teams.

Jensen's [article](#) in *Information, Communication & Society*, co-authored with **Samuel Baker** (English, College of Liberal Arts) and **Dhiraj Murthy** (School of Journalism and Media, Moody College of Communication), examines how generative AI reconfigures the relationship between creativity, labor and power. Building on that work, Jensen completed an invited chapter to the forthcoming edited volume "Teaming with Machines: How Contemporary Teams Collaborate with Complex Work Technologies." Additionally, in a new study of how generative AI is transforming the music industry, he built a database of generative AI tools and initiated interviews with music industry professionals.

In a second line of research, Jensen examined how interdisciplinary teams negotiate ethical responsibility in AI research. Drawing on interviews of interdisciplinary research teams, the study theorizes “ethical boundary coordination” — the process through which researchers with different disciplinary commitments define, negotiate and translate shared ethical responsibilities in practice. Jensen and Murthy presented the work at the 2026 International Communication Association conference in Cape Town, South Africa.

Jensen has accepted a tenure-track position at the University of San Diego beginning in fall 2026.

[READ MORE](#)

Seed Projects

Good Systems awarded seed funding to six interdisciplinary research projects selected from a pool of 48 proposals. Launched in January 2026, the projects extend the grand challenge’s network of researchers into new domains, from maternal health and flood resilience to disability access and K-12 education.

ADDRESSING MATERNITY CARE DESERTS

Yuhao Kang (Geography and the Environment, College of Liberal Arts) and **Lorie Harper** (Women’s Health, Dell Medical School) lead [this project](#), which is building an AI- and GIS-based decision-support tool to identify gaps in maternity care access across Texas. **Hyeun Ah Kang** (Health Outcomes, College of Pharmacy) and **Rajesh Reddy** (Maternal-Fetal Medicine, Dell Medical School) are also investigators on the project.

AI²: ANCESTRAL INTELLIGENCE AND ARTIFICIAL INTELLIGENCE FOR FLOOD MITIGATION

Lidia Cano Pecharromán (School of Architecture) and **ChangHoon Hahn** (Astronomy, College of Natural Sciences) are redesigning an [AI-driven flood-evaluation tool](#) so tribal nations can incorporate Indigenous worldviews and community-focused measures of effectiveness into their flood mitigation programs.

AUTONOMOUS VEHICLES AND ARTIFICIAL INTELLIGENCE FOR ALL (AVAIL)

Zhaomiao Guo (Maseeh Department of Civil, Architectural and Environmental Engineering, Cockrell School of Engineering) and his team — **Wen (Vivian) Ye** (School of Design and Creative Technologies, College of Fine Arts) and **Qixing Huang** (Computer Science, College of Natural Sciences) — are developing an AI-enabled mobility assistance system to improve safety and independence for blind and visually impaired people interacting with autonomous vehicles.

HACK THE CAMP-US

Lillian Chin (Chandra Family Department of Electrical and Computer Engineering, Cockrell School of Engineering) leads [this project](#), which uses community-driven design to create AI-enabled interventions that address infrastructural barriers for people with disabilities at UT Austin. The co-PIs represent a wide range of disciplines and colleges: **Patrick Benfield** (College of Natural Sciences); **Leah Chong** (Walker Department of Mechanical Engineering, Cockrell School of Engineering); **Jo Hsu** (Rhetoric and Writing, College of Liberal Arts); **DiMitri Higginbotham** (School of Design and Creative Technologies, College of Fine Arts); **Earl Huff Jr.** (School of Information); **Alison Kafer** (Women's, Gender and Sexuality Studies, College of Liberal Arts); and **Sandy Magaña**, **Angela Standridge** and **Emily Shryock** (all from the School of Social Work).

KNOWLEDGE-INFORMED MULTIMODAL RESPONSIBLE AI FOR COPD

Edward Castillo (Biomedical Engineering, Cockrell School of Engineering) leads this team developing an AI framework to improve chronic obstructive pulmonary disease care while making the model's outputs easier for clinicians to interpret, so clinical decisions are more transparent and patient centered. **Hairong Wang** (Walker Department of Mechanical Engineering, Cockrell School of Engineering) and **Ying Ding** (School of Information) are co-PIs.

YOUTH PERSPECTIVES ON TEXT-TO-IMAGE AI

Angela Smith (School of Information) and **Patricia Abril-Gonzalez** (Curriculum and Instruction, College of Education) are working with K-12 students to generate youth-driven recommendations for more ethical and equitable generative AI.

[READ MORE](#)



CREATING CONNECTIONS

Good Systems' greatest asset is its network. Faculty and students from nearly every college and school at UT Austin take part in the grand challenge, and that breadth is what makes its interdisciplinary, cross-sector work possible. The network continued to grow in FY26, as a new cohort of seed grants drew researchers from fields as varied as geography, astronomy, biomedical engineering and education.

Beyond the University, Good Systems deepened established relationships and forged new ones. Locally, it sustained its longstanding partnership with the City of Austin, remained a founding member of the Austin AI Alliance and an active voice in the Austin Forum on Technology & Society, and extended its reach through Austin Community College and Huston-Tillotson University, where researchers took part in the 2026 HBCU AI Conference. Nationally and internationally, the Alliance for Interdisciplinary AI — the community of practice Good Systems founded among peer institutions — continued to grow, and the grand challenge remained engaged with partners such as the Public Interest Technology University Network (PIT-UN) and Responsible AI UK.

Within the Forty Acres, a new partnership with Enterprise Technology strengthened AI education and research on the responsible adoption of AI technologies across campus, while continued collaboration with the Machine Learning Lab and Texas Robotics expanded the community around responsible AI and robotics at UT. Ties with industry strengthened as well, most notably with AMD, whose responsible AI team became a closer collaborator over the year.

Good Systems also remained a trusted resource for policymakers, contributing to AI governance discussions. This included helping shape the Texas Responsible Artificial Intelligence Governance Act, which was signed into law by Governor Greg Abbott in 2025.

Explore Good Systems' relationships through our [interactive network map](#). Search by name, by college, school or unit, or by project; select any node to view its connections; and zoom in to trace how a single researcher has helped shape the grand challenge's work.

Sharing the Story of Good Systems

In FY26, Good Systems released a series of short films capturing the grand challenge's research and its stakes. The centerpiece is a [six-minute film](#) on the research impact of the grand challenge and the continuing need for ethical AI research. It is complemented by several shorter videos, including one about how Good Systems researchers [define ethical AI](#) and another about the questions to ask when [evaluating](#) the ethics of AI systems. The series draws on interviews with Good Systems researchers and leaders, as well as external partners representing Responsible AI UK, the Carnegie Endowment for International Peace, KUNGFU.AI and Austin City Council. Designed to reach a broad audience, the videos aim to strengthen partnerships, foster public trust and encourage further investment in interdisciplinary AI ethics research.



Research Engagement

Because responsible AI research affects people far beyond the University, Good Systems makes public engagement part of the work. Its free public programs bring researchers, practitioners and community members together to examine the promises and risks of AI while strengthening the partnerships that shape the grand challenge's research.

TEXAS SYMPOSIUM ON MACHINE LEARNING, RESPONSIBLE AI, AND ROBOTICS

In partnership with the Machine Learning Lab and Texas Robotics, Good Systems held the inaugural Texas Symposium on Machine Learning, Responsible AI and Robotics in March 2026, expanding its annual symposium into a larger campuswide convening. The program brought together the research and partner networks of UT's machine learning, robotics and responsible AI communities, creating new opportunities for exchange across

the technical and societal dimensions of AI research. “Right now we’re in a moment where machine learning and artificial intelligence — especially generative artificial intelligence — are changing the world in many ways,” Peter Stone (Computer Science, College of Natural Sciences) said. “This is the first time these three organizations have come together to give a joint symposium, but I think, especially in this moment, it’s fitting to have machine learning, robotics and Good Systems all together in this room.”

With more than 650 attendees, the event showcased the scale of UT Austin’s AI research community and created a shared forum for discussing how AI systems are designed, deployed and governed.

[READ MORE](#)

RESPONSIBLE AI IN 2025

In November, Good Systems brought together researchers and AI leaders from around the country for Responsible AI in 2025: State of the Art and Future Directions, a two-day online symposium. In panels led by each core research team, speakers reflected on progress and open challenges in developing responsible AI technologies for society and shared research-based insights related to work, robotics, smart cities, the information environment and more. The conversations highlighted the importance of interdisciplinary and cross-sector collaboration in shaping accountable and human-centered AI systems.

[READ MORE](#)

“This is the first time these three organizations have come together to give a joint symposium, but I think, especially in this moment, it’s fitting to have machine learning, robotics and Good Systems all together in this room.”

PETER STONE, COLLEGE OF NATURAL SCIENCES

RESPONSIBLE AI FOR EDUCATION

In April 2026, Good Systems co-hosted the second annual Responsible AI for Education event with the Office of Academic Technology and Enterprise Technology. The event helped connect Good Systems' responsible AI expertise with campuswide conversations about teaching, learning and technology adoption. Keynotes, panels and workshops explored student, faculty and staff perspectives on AI use in research and instruction, with recurring attention to critical thinking, sustainability, privacy and the value of using AI to support rather than replace learning. **James P. Frazee**, chief information officer (CIO) at San Diego State University, joined UT Austin CIO **Cole Campese** to discuss larger trends in technology adoption and how to build a stakeholder-centered approach to technology deployment on university campuses.

[READ MORE](#)



Good Systems co-hosted the second annual Responsible AI for Education symposium on campus on April 21.

SMART CITIES AND AI INNOVATIONS SYMPOSIUM

The fourth annual Smart Cities and AI Innovations Symposium brought together researchers, practitioners and leaders from academia, industry, government and nonprofits to examine how AI is changing urban life. Organized by researchers from A Good System for Smart Cities, the program connected traditional smart-city concerns, such as infrastructure, transportation and housing, with emerging areas like accessibility, health care and workforce development. Speakers emphasized that responsible urban AI depends on cross-sector collaboration, strong and equitable data infrastructure, and close attention to local needs. The symposium also highlighted next-generation urban AI research at UT Austin, with postdoctoral fellows and doctoral students presenting work on digital twins, GeoAI, flood modeling and LLM accountability.

[READ MORE](#)



On June 4, researchers, practitioners and leaders from academia, industry, government and nonprofits came together on campus for the fourth annual Smart Cities and AI Innovations Symposium, organized by leaders from A Good System for Smart Cities, one of Good Systems' core research projects.

Collaborations with Academic, Community and Industry Partners

Throughout FY26, Good Systems deepened its collaborations across the University, with Austin community and industry leaders, and with academic institutions nationwide.

ETHICS UNWRAPPED

Good Systems contributed to **“Running with Scissors: AI and the Race for the Future,”** a documentary produced by Ethics Unwrapped at the McCombs School of Business as part of its AI Ethics Docuseries. The film examines how the rapid development and deployment of AI can outpace ethical reflection and features Good Systems faculty, students and partners whose work has helped define the grand challenge’s approach to responsible AI.

The collaboration extended Good Systems’ research into a widely used public teaching resource. Ethics Unwrapped reaches audiences well beyond UT Austin, with materials used on more than 1,800 campuses and in businesses and organizations around the world, making the documentary a useful vehicle for bringing Good Systems perspectives on AI ethics, accountability and human-centered design to students, educators and professionals.

The following researchers participated in the Ethics Unwrapped documentary:

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

Maria De-Arteaga, ESADE Business School

Emily Economy, graduate student, Communication Studies, Moody College of Communication

Kenneth Fleischmann, School of Information

Sherri Greenberg, LBJ School of Public Affairs

Steve Kramer, KUNGFU.AI

Tina Lassiter, graduate student, School of Information

Matthew Lease, School of Information

Chelsea McCullough, former graduate student, School of Information

Keri Stephens, Communication Studies, Moody College of Communication

Peter Stone, Computer Science, College of Natural Sciences

Sharon Strover, School of Journalism and Media, Moody College of Communication

Yifan Xu, former graduate student, Communication Studies, Moody College of Communication

ENTERPRISE TECHNOLOGY

Good Systems partnered with UT Austin's Enterprise Technology team to advance AI education and to study responsible AI adoption on campus and the needs and concerns of users. One line of inquiry, led by doctoral student **Jaxsen Day** with support from **Kenneth Fleischmann** (both from the School of Information) and collaborator **Mario Guerra** (Enterprise Technology), examines how AI tools can help students with visual disabilities find, navigate and understand research materials that are not readily accessible.

[READ MORE](#)

Enterprise Technology also supported research led by **Min Kyung Lee** (School of Information) and **Sherri Greenberg** (LBJ School of Public Affairs) on responsible generative AI governance and privacy policy in organizations. **Sally You**, a doctoral student in the School of Information, helped launch a survey to better understand how students, faculty and staff think about the benefits and perceived privacy risks of using generative AI in higher education settings.

ALLIANCE FOR INTERDISCIPLINARY AI

The Alliance for Interdisciplinary AI — the community of practice for AI ethics researchers across the U.S. co-founded by Good Systems — continued to meet quarterly and held its annual in-person meeting at the University of Oklahoma in May 2026. During the meeting, participants shared insights on issues in responsible AI research and teaching at U.S. universities, from the governance of agentic AI to creative approaches to teaching with AI and the changing research landscape. New this year, the community of practice hosted a public-facing panel focused on the role of universities in the age of AI and featured several members as expert speakers, including **Sharon Strover** (Journalism and Media, Moody College of Communication).

The group has grown to 34 researchers from 12 universities. This year its members formed several working groups to advance shared goals, including creating a shared teaching-resources hub and identifying new opportunities for collaborative research, publications and public programs, and they established an official mission statement, vision statement and brand identity.

LOCAL AND COMMUNITY PARTNERS

In FY26, Good Systems strengthened its local and community ties across Austin. The grand challenge remained a founding member of the Austin AI Alliance and a member of the Austin Forum on Technology & Society. It also engaged with Austin Community College and Huston-Tillotson University, where researchers took part in the HBCU AI Conference in March 2026.

The grand challenge also deepened its partnership with AMD, which hired a Good Systems research assistant, mechanical engineering doctoral student **Xingru Zhou**, to study agentic AI in the Human-Centered Robotics Lab led by **Luis Sentis** (Aerospace Engineering and Engineering Mechanics, Cockrell School of Engineering). AMD's global lead for responsible AI, Aarti Choudhary, participated in a **panel** at the Texas Symposium on Machine Learning, Responsible AI and Robotics.



AMD's global lead for responsible AI, Aarti Choudhary (right), was a panelist at the inaugural Texas Symposium on Machine Learning, Responsible AI and Robotics, on March 4, 2026. Elias Stengel-Eskin (Computer Science, College of Natural Sciences) and Lucas Meyer (Principal Research Scientist, Microsoft AI for Good Lab) joined her on the panel, "Reimagining Work in the Age of AI Agents." Photo credit: Noah Devereaux

Good Systems Partners

Good Systems collaborates with external organizations and institutions across industry, the nonprofit sector, government and academia.

INDUSTRY

- AMD
- The Cantellus Group
- Full Fact
- Good Tech Advisory LLC.
- KUNGFU.AI
- Meedan
- Smart Cities Connect

NONPROFITS

- Austin AI Alliance –
Good Systems is a founding member
- Austin Forum on Technology & Society –
Good Systems is a member
- Data & Society
- MEASURE
- MITRE
- Public Interest Technology
University Network (PIT-UN)

GOVERNMENT

- Austin City Council
- City of Austin
- City of San Antonio
- NIST AI Consortium –
Good Systems is a member
- Texas Department of Information
Resources (DIR)
- Texas Innovation &
Technology Caucus
- Texas Workforce Commission

ACADEMIC INSTITUTIONS

- Austin Community College
- Carnegie Mellon University
- Huston-Tillotson University
- Johns Hopkins University Institute
for Assured Autonomy
- NC State University Center for
AI in Society & Ethics
- Penn State Center for Socially Responsible AI
- Responsible AI UK
- Rutgers University Critical AI
- UKRI Trustworthy Autonomous
Systems Hub
- University of Buffalo
- University of Cincinnati
- University of Georgia
- University of Maryland
- University of Oklahoma
- University of Richmond
- University of Washington

GOOD SYSTEMS HEADLINES

EXTERNAL FEATURES AND MENTIONS

- 10/8/2025 [Artificial intelligence expert weighs in on fake home invasion TikTok prank](#)
FOX 7
- 2/11/2026 [How Can AI Companions Be Helpful, not Harmful?](#)
IEEE Spectrum
- 3/12/2026 [The Anthropic Pentagon Standoff and the Limits of Corporate Ethics](#)
Tech Policy Press
- 4/1/2026 [From broken PDFs to instant access: How ChatGPT rebuilds the research workflow at UT Austin](#)
OpenAI Academy
- 4/8/2026 [Studying the Risks and Benefits of AI Companions: Researchers Discuss a New Framework for Understanding AI Companionship](#)
Sony AI
- 4/29/2026 [Austin's growing AI scene fuels debate: innovation or regulation?](#)
KXAN
- 6/17/2026 [Do people feel safe in a robot's presence?](#) Science
Robotics
- 7/1/2026 [UT Austin team examining ethics of AI use to release 'People's Guide to AI' in 2026](#)
KXAN

UT NEWS COVERAGE

- 9/29/2025 [AI to the rescue? UT research team develops multilingual emergency messaging bot](#)
Texas Connect
- 11/5/2025 [From Research to Results: UT-City of Austin Partnership Delivers Regional Innovation](#)
UT News
- 2/17/2026 [Seed Funding for Human-Centered AI Projects Awarded to Three COLA Faculty](#)
College of Liberal Arts
- 3/20/2026 [Leaders in AI, Robotics and Ethical Innovation Come Together at UT Austin](#)
UT News
- 3/26/2026 [Interdisciplinary faculty team advances human-centered AI to empower tribal nations in flood planning](#)
School of Architecture
- 6/25/2026 [iSchool Researchers Bring Open-Source AI, Fairness, and Accountability to ACM FAccT 2026](#)
School of Information
- 8/27/2026 [UT To Lead NSF Center That Will Study How Robots and People Learn To Live and Work Together](#)
UT News

UT OVPR COMMUNICATIONS

- 11/24/2025 [Cross-Cutting Edge](#)
- 12/17/2025 [Meaning, Making and Machine Learning](#)
- 1/26/2026 [Good Systems Awards Seed Grants to Six Faculty Teams Advancing Human-Centered AI](#)
- 5/20/2026 [Ethics First: How Good Systems is Helping the City of Austin Navigate the AI Frontier](#)
- 8/26/2026 ['We Chose the Right Problem'](#)

FUNDED GRANTS

To date, Good Systems researchers have secured more than \$52,775,000 in external funding through grants, gifts and awards that directly support the grand challenge's mission. External awards from the most recent fiscal year are listed below.

Amazon Research Award

From Observation to Intervention:

Counterfactual Multi-Agent World Models for Autonomous Driving

\$60,000

Zhaomiao Guo, Cockrell School of Engineering

Honda Research Institute

Learning from Implicit Human Feedback for Adaptive In-Vehicle Systems

\$501,951

Brad Knox, College of Natural Sciences

IBM

Governance and Training Frameworks for Ethical Generative

AI Adoption in City Governments

\$20,000

Sharon Strover, Moody College of Communication

National Science Foundation

Science and Technology Center for Human and Robot Co-Adaptation

\$30,000,000

Joydeep Biswas, College of Natural Sciences

***U.S. Department of Defense**

Risks from AI-based Companionship: An Ontological, Causal Account of Potential Harms

\$60,000

Brad Knox, College of Natural Sciences

*Awarded in FY25 but not previously reported

PUBLICATIONS

During FY26, Good Systems researchers advanced knowledge through 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, all of which are available online.

Agarwal, K., Jiang, Y., Hu, J., Liu, B., & Stone, P. (2025, October). **L3M+P: Lifelong planning with large language models** [Paper presentation]. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Hangzhou, China.

Angliss, C., Cui, J., Hu, J., Rahman, A., & Stone, P. (2026, May). **VGC-bench: Towards mastering diverse team strategies in competitive Pokémon** [Paper presentation]. International Conference on Autonomous Agents and Multiagent Systems, Paphos, Cyprus.

[†]Arceneaux, P., Anderson, J., Lukito, J., Shah, M., & Kioussis, S. (2026). **Social bots as agenda-builders: Evaluating the impact of algorithmic amplification on organizational messaging**. *Journal of Public Relations Research*, 38(3), 123–156.

Beliveau, A. E., Haddad, A. J., Podnar, E. A., Sharma, D., & Bhat, C. R. (2026). **Disparities in pedestrian crossing and driver yielding behaviors: Evidence from a large-scale observational study at urban intersections**. *Accident Analysis & Prevention*, 230, 108465.

[†]Brown, M., Lukito, J., Macdonald, M., Dowling, K., Hickey, C., & Miranda, M. (2026, May). **Candidata: U.S. 2024 elections candidates and social media posts** [Paper presentation]. AAAI Conference on Web and Social Media, Los Angeles, CA.

Chang, A., & Jiao, J. (2026). **Energy creates intelligence: The infrastructural geography of global AI**. *AI & Society*, 41, 6781-6795.

[†]Dahlke, R., Yang, Y., Lukito, J., Greenfield, J., Chen, B., Brown, M. A., & Lewis, R. (2026). **Style and substance on The Alex Jones Show predict InfoWars sales: A multi-modal analysis of a media empire**. *Information, Communication & Society*, 29(4), 1386–1405.

Gautam, S., Liu, H., Choi, Y., & Lease, M. (2026, June). **How researchers navigate accountability, transparency, and trust when using AI tools in early-stage research: A think-aloud study** [Paper presentation]. ACM Conference on Fairness, Accountability, and Transparency (FAcCT), Montréal, Canada.

*Gupta, R., Norman, E., Shin, H., Deng, Z., Esteva, M., Lu, N., Stephens, K. K., & Sentis, L. (2026). **Indoor human-mobile robot encounters: A transdisciplinary study on perceived safety**. *ACM Transactions on Human-Robot Interaction*, 15(4), 1-41

*Gupta, S., De-Arteaga, M., & Lease, M. (2026, June). **Fairness-aware multi-group target detection in online discussion** [Paper presentation]. ACM Conference on Fairness, Accountability, and Transparency (FAcCT), Montréal, Canada.

Handman, C. (2026). **Animal translations: AI and the intelligibility of non-human worlds**. *Journal of the Royal Anthropological Institute*, 32, 1003-1021.

- [†]Hila, A., & Hauser, E. (2025). **Assessing the reliability of large language models for deductive qualitative coding: A comparative intervention study with ChatGPT**. *Proceedings of the Association for Information Science and Technology*, 62(1), 275–285.
- Hu, J., Stone, P., & Martín-Martín, R. (2025, September). **SLAC: Simulation-pretrained latent action space for whole-body real-world RL** [Paper presentation]. Conference on Robot Learning, Seoul, Korea.
- *Jensen, J. T., Murthy, D., & Baker, S. (2025). **Automating remix: Generative AI, creative labor, and the decay of aura**. *Information, Communication & Society*, 29(12), 3603–3619.
- Jia, C., Gondimalla, A., Zhang, A., Mullings, D. J., Boltz, A., & Lee, M. K. (2026). **Beyond accuracy: Experts see AI fact-checks as accurate but less useful**. *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*, 929, 1–22.
- Jiao, J., Afroogh, S., Chen, K., Atkinson, D., & Dhurandhar, A. (2026). **Generative AI and LLMs in industry: A text-mining analysis and critical evaluation of guidelines and policy statements across 14 industrial sectors**. *Humanities and Social Sciences Communications*, 13(1), 410.
- *Jiao, J., Afroogh, S., Murali, A., Chen, K., Atkinson, D., & Dhurandhar, A. (2025). **LLM ethics benchmark: A three-dimensional assessment system for evaluating moral reasoning in large language models**. *Scientific Reports*, 15(1), 34642.
- Jiao, J., Afroogh, S., Xu, Y., & Phillips, C. (2025). **Navigating LLM ethics: Advancements, challenges, and future directions**. *AI and Ethics*, 5, 5795–5819.
- *Joseph, S., Chen, L., Wei, B., Mackert, M., Marshall, I. J., Liang, P. P., Kouzy, R., Wallace, B. C., & Li, J. J. (2026). **Decide less, communicate more: On the construct validity of end-to-end fact-checking in medicine**. *Findings of the Association for Computational Linguistics*, 10196–10221.
- *Kambhatla, G., Gautam, S., Zhang, A., Liu, A., Srinivasan, R., Li, J. J., & Lease, M. (2026, July). **Improving the distributional alignment of LLMs using supervision** [Paper presentation]. The Association for Computational Linguistics, San Diego, CA.
- *McCullough, C., Needle, J. P., Triem, H., Fleischmann, K. R., & Greenberg, S. R. (2025). **“AI is not gonna take our jobs”: Perspectives from skilled trade labor union apprentices**. *IEEE Transactions on Technology and Society*, 6(4), 471–480.
- Mo, K., Venkatayogi, S., Shaib, C., Kouzy, R., Xu, W., Wallace, B. C., & Li, J. J. (2026, July). **Faithfulness vs. safety: Evaluating LLM behavior under counterfactual medical evidence** [Paper presentation]. The Association for Computational Linguistics, San Diego, CA.
- Munje, M. J., Tang, C., Liu, S., Hu, Z., Zhu, Y., Cui, J., Warnell, G., Biswas, J., & Stone, P. (2025, September). **SocialNav-SUB: Benchmarking VLMs for scene understanding in social robot navigation** [Paper presentation]. Conference on Robot Learning, Seoul, Korea.
- New, R., & Longoria, R. G. (2025). **Evaluating skill detection methods for smart powered hand tools**. *Conference Proceedings of the 2025 International Design Engineering Technical Conferences & Computers and Information in Engineering, Volume 2A*, V02AT02A053.
- *[†]Saab, M., Chang, A., Baik, S. G., Jiao, J., & Rausch, C. (2026). **Digital twins for accessory dwelling unit feasibility under environmental and infrastructure constraints: The case of Austin, Texas**. *Applied Spatial Analysis and Policy*, 19, 151.

- *Salazar, C. D., New, R., McCullough, C., Lassiter, T., Fleischmann, K. R., Greenberg, S. R., Sha, Z., & Longoria, R. G. (2025). **Design, development, and testing of smart hand tool systems**. *Conference Proceedings of the 2025 International Design Engineering Technical Conferences & Computers and Information in Engineering, Volume 2A*, V02AT02A058.
- *Seong, K., Jiao, J., Lewis, R. H., Farahi, A., Navrátil, P., Casebeer, N., Davis, B., Jones, J., & Niyogi, D. (2025). **Towards a digital twin for smart resilient cities: Real-time fire and smoke tracking and prediction platform for community awareness (FireCom)**. *Computational Urban Science*, 5(1), 49.
- Shensky, M. G., Strickland, K. M., Marden, A. W., & Dubbe, H. (2024). **An automated georeferencing workflow for processing historical sanborn fire insurance maps using object detection and optical character recognition**. *Journal of Map & Geography Libraries*, 20(3), 137–163.
- Su, Y., & Lease, M. (2026, October). **LLM REgression with a latent iterative state head**. [Paper presentation]. Conference on Language Modeling, San Francisco, CA.
- *Wang, H., Davis, W., Xu, Y., Yu, J., Mai, G., & Jiao, J. (2025). **Street semantic tree: A knowledge-driven GeoAI framework for urban e-scooter ridership classification**. *International Journal of Geographical Information Science*, 40(6), 1680-1709.
- Xu, Y., & Jiao, J. (2026, January). **Evaluating retrieval-augmented generation strategies for large language models in travel mode choice prediction** [Paper presentation]. Transportation Research Board 105th Annual Meeting, Washington D.C.
- *Xu, Z., Procko, K., Munje, M., Patterson, K., Sabatini, L., Biswas, J., & Stone, P. (2026, January). **The essentials of AI for life and society: A full-scale AI literacy course accessible to all**. Symposium on Educational Advances in Artificial Intelligence (EAAI). [Paper presentation]. Singapore, Singapore.
- Yang, H., Lee, J.-W., & Santos Cruz, A. U. (2026). **Machine learning-enhanced static flood models for high-resolution peak storm surge inundation mapping in Southeast Texas, USA**. *Journal of Hydrology: Regional Studies*, 63, 103056.
- *Zhang, A., Liao, M., Kravchenko, E. (Lee), 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*, 9(7), 1–27.
- Zhang, A., & Lee, M. K. (2026). **Contextualizing datasets: Deepening awareness of data biases through critical reflection of civic data**. *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*, 413, 1–29.
- Zhang, X., Ding, Y., Hayamizu, Y., Altaweel, Z., Zhu, Y., Zhu, Y., Stone, P., Paxton, C., & Zhang, S. (2025). **LLM-GROP: Visually grounded robot task and motion planning with large language models**. *The International Journal of Robotics Research*, 0(0), 02783649251378196.
- Zheng, A. A., Li, J. J., & Beaver, D. I. (2026). **Strategic dialogue assessment: The crooked path to innocence**. *Dialogue & Discourse*, 17(1), 1–53.

* Denotes scholarly outputs where there are authors from more than one UT department or school.

† Denotes publications that DO NOT formally acknowledge Good Systems support; all others do.

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.

Kenneth R. Fleischmann

Chair (2025-26) and
Founding Chair
School of Information

Luis Sentis

Past 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

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

Stacey Ingram Kaleh

Associate Director,
Partnerships and Programs
Office of the Vice President for
Research

Core Research Projects

DESIGNING AI FOR SOCIAL FAIRNESS

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

Min Kyung Lee

School of Information
Project Co-Lead

Kaya de Barbaro

Psychology
College of Liberal Arts

Hwarang Stephen Han

Medicine
Dell Medical School

Lorie Harper

Women's Health
Dell Medical School

Matt Kammer-Kerwick

IC² Institute

Alex Karner

School of Architecture

Lisa Macias

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

Ruben Rathnasingham

Medical Education
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

Joydeep Ghosh

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

Edgar Gómez-Cruz

School of Information

Jianing Zhu

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

SMART HAND TOOLS: BUILDING THE FUTURE OF WORK WITH ETHICAL AI

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

Sanjana Gautam

School of Information

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-Tsechansky

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

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

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

Aaron Choate

University of Texas Libraries

Maria Esteva

Texas Advanced Computing
Center

Courtney Handman

Anthropology
College of Liberal Arts

Eliana Hadjiandreou

Psychology
College of Liberal Arts

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

Keri Stephens

Communication Studies
Moody College of Communication

Peter Stone

Computer Science
College of Natural Sciences

Nathan Tsoi

Computer Science
College of Natural Sciences

Sean Williams

School of Law

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

Communication Studies
Moody College of Communication

Seed-Funded Projects

ADDRESSING MATERNITY CARE DESERTS: QUANTIFYING MULTI-LEVEL HEALTHCARE ACCESS TO OPTIMIZE RESOURCE ALLOCATION WITH AI

Lorie Harper
Department of Women’s Health
Dell Medical School

Hyeun Ah Kang
College of Pharmacy

Yuhao Kang
Geography and the Environment
College of Liberal Arts

Rajesh Reddy
Medical Education
Dell Medical School

AI^2: ANCESTRAL INTELLIGENCE AND ARTIFICIAL INTELLIGENCE FOR FLOOD MITIGATION

ChangHoon Hahn
Astronomy
College of Natural Sciences

Lidia Cano Pecharromán
School of Architecture

AUTONOMOUS VEHICLES AND ARTIFICIAL INTELLIGENCE FOR ALL (AVAIL)

Zhaomiao Guo
Maseeh Department of Civil,
Architectural and Environmental
Engineering
Cockrell School of Engineering

Qixing Huang
Department of Computer Science
College of Natural Sciences

Wen (Vivian) Ye
School of Design and Creative
Technologies
College of Fine Arts

HACK THE CAMP-US: SOCIOTECHNICAL INTERVENTIONS BY AND FOR PEOPLE WITH DISABILITIES

Patrick Benfield

Office of Experiential Learning
College of Natural Sciences

Lillian Chin

Electrical and Computer
Engineering
Cockrell School of Engineering

Leah Chong

Walker Department of
Mechanical Engineering
Cockrell School of Engineering

DiMitre Higginbotham

School of Design and
Creative Technologies
College of Fine Arts

Jo Hsu

Rhetoric and Writing
College of Liberal Arts

Earl Huff Jr.

School of Information

Alison Kafer

Women's, Gender and
Sexuality Studies
College of Liberal Arts

Sandy Magaña

School of Social Work

Emily Shryock

Disability Cultural Center
School of Social Work

Angela Standridge

Texas Technology
Access Program
School of Social Work

KNOWLEDGE-INFORMED MULTIMODAL RESPONSIBLE AI FOR COPD

Edward Castillo

Biomedical Engineering
Cockrell School of Engineering

Ying Ding

School of Information

Hairong Wang

Walker Department of
Mechanical Engineering
Cockrell School of Engineering

YOUTH PERSPECTIVES ON TEXT-TO-IMAGE AI: CO-DESIGNING ETHICAL AND EQUITABLE GENERATIVE TECHNOLOGIES FOR K-12

Patricia Abril-Gonzalez

Curriculum and Instruction
College of Education

Angela Smith

School of Information