Curriculum Vitae

Pradeep Kumar Murukannaiah

Van Mourik Broekmanweg 6 – 2628 XE Delft – The Netherlands

+31 06 51937995 • p.k.murukannaiah@tudelft.nl
https://ii.tudelft.nl/~pradeep/

Updated: October 4, 2024

Contents

Executive Summary

Background: I am an Assistant Professor in the Intelligent Systems department at TU Delft. I received a PhD and MS in Computer Science from North Carolina State University (Raleigh, USA) and a BE in Information Science from University Visvesvaraya College of Engineering (Bangalore, India). Prior to TU Delft, I was an Assistant Professor at Rochester Institute of Technology (Rochester, USA).

Research: My research aims to connect **dialogue** and **decisions** on **societal issues** such as climate change. I pursue this goal via three research lines. I investigate (1) Natural Language Processing (NLP) techniques to identify what stakeholders *value* in a societal dialogue, and (2) Multi-Objective Optimization (MOO) techniques to analyze the stakeholders' values (extracted via NLP), and (3) Hybrid Intelligence (HI) methods, where the artificial intelligence (NLP and MOO) augments human intelligence for societal decision making. I strive to publish in high-quality venues and my works have appeared in 27 top-tier (A*/A) Computer Science conferences and journals. My works have been recognized with a Best-Paper award, a Best-BlueSky Paper award, and a Best Poster award at international conferences.

Valorization: My research topics are **scientifically novel** (e.g., values in NLP) and **societally significant** (e.g., MOO decision-support on climate change). As a result, I have been able to acquire funding to support two post-doctoral associates, six PhD students, and a research engineer. The sources of my funding include: two NWO grants (two PostDocs and one research engineer), Hippo Delft AI lab, supported by TU Delft AI Initiative (two PhDs), the Hybrid Intelligence Center, funded by a 10 year Zwaartekracht grant (three PhDs), and BOLD Cities. In addition, I have a network of strong collaborators in TU Delft (e.g., in EWI, TPM, and ABE), nationally (e.g., in VU Amsterdam, Utrecht, Leiden, RU Groningen, Erasmus, Tilburg, and THUAS), and internationally (e.g., in US, UK, and Spain).

Teaching: I teach **foundational** (e.g., Computational Intelligence), **state-of-the-art** research-driven (e.g., Collaborative AI), **interdisciplinary** (e.g., Advanced Interdisciplinary AI Project) as well as **societally significant** (e.g., NLP for Society) courses. Further, I have experience in multiple facets of teaching, including **lecturing** (three courses), **coordination** (three courses), **development** (three courses) and **supervision** of MSc theses (seven in Computer Science and one each in TPM and ABE) and BSc projects (30 in Computer Science). At Rochester Institute of Technology (my previous institute), I led an REU (Research Experiences for Undergraduates) site on Software Engineering, funded by the US National Science Foundation (USD 360,000). This REU site provided research experience to undergraduate students, especially, those from universities that did not have adequate research opportunities.

Leadership: I currently **lead a research team** of one PostDoc, four PhD students, and three MSc students. My first priority, as a leader, is to ensure the success of my team. All students I have previously supervised have successfully graduated within the expected time frame, and my current students are all on track to successfully complete their degrees on time. Beyond my immediate team, I am entrusted with important roles in **leading research consortia**. For example, I led one of the four research lines of the Hybrid Intelligence consortium, which consists of over 100 researchers. I made an impact in this role not only by developing a research agenda but also by fostering collaborations. My contributions to the consortium are acknowledged by the additional key responsibilities I have been entrusted with, e.g., leading four use cases and representing the consortium for mid-term project renewal. Even broadly, I **lead the scientific community** both at a content level (e.g., by publishing vision articles and editing special issues) and at an organizational level (e.g., by being Program or Track Chair for key conferences).

Organization: I have been playing three key organizational roles at TU Delft. First, in a working group of three, I developed two proposals on **AI education for all TU Delft students**, contributing to the TU Delft AI Initiative's vision. These proposals led to the creation of two 15 EC MSc courses, Foundations of AI Program (FAIP) and Advanced Interdisciplinary AI Project (AI2P). I also serve as the course coordinator for AI2P. Second, I was a member of a **core team** that led the development of the new **MSc DSAIT** program. As part of this team, I investigated similar programs, developed the intended learning outcomes, and established a coherent program structure. Recognizing my contribution, I was chosen as a program representatives for the NVAO panel, and subsequently, I was appointed as one of the two Master Coordinators for the program. Finally, I am serving in the management team of the **Delft AI Labs and Talent** program. Our team develops strategy and oversees the research, innovation, education, and outreach activities of the 24 AI labs and several AI talents of TU Delft.

Professional Background

Education

- **PhD in Computer Science** — May 2016
  North Carolina State University, Raleigh, USA
  Outstanding Dissertation in Computer Science Award
- **MS in Computer Science** — Dec 2011
  North Carolina State University, Raleigh, USA
  GPA: 4.0/4.0
- **BE in Information Science and Engineering** — May 2005
  University Visvesvaraya College of Engineering, Bangalore, India
  First Class with Distinction

Training and Certification

- **University Teaching Qualification (UTQ)** — Sep 2022
  Delft University of Technology, Delft, The Netherlands
  Modules: SUPERVISE, DESIGN, TEACH, ASSESS (Dossier)
- **English Language Test for Lecturers** — May 2024
  Delft University of Technology, Delft, The Netherlands
  C2 (highest) Proficiency

Academic Experience

- **Assistant Professor 1** — Mar 2020–Present
  Assistant Professor 2 — Aug 2019–Mar 2020
  Interactive Intelligence, Intelligent Systems
  Faculty of Electrical Engineering, Mathematics, and Computer Science (EEMCS),
  Delft University of Technology, Delft, The Netherlands
- **Assistant Professor** — Jan 2017–Aug 2019
  Department of Software Engineering, Center for Cybersecurity (Affiliated Faculty),
  B. Thomas Golisano College of Computing and Information Sciences (GCCIS),
  Rochester Institute of Technology, Rochester, NY, USA
- **Postdoctoral Research Scientist** — May 2016–Jan 2017
  Department of Software Engineering, GCCIS,
  Rochester Institute of Technology, Rochester, NY, USA
- **Data Intern** — May 2009–Aug 2009
  Neurobiology of Vocal Communication Lab (Transcriptome Data Migration)
  Duke University Medical Center, Durham, NC, USA

Industry Experience

- **Software Engineering Intern** — May 2011–Aug 2011
  AdPlanner Demographics Team (Machine learning),
  Google, Inc., Seattle, USA
- **Software Engineer** — Jul 2005–Dec 2008
  Optical Networks Management System (Database admin. and Java development),
  Alcatel-Lucent India Pvt. Ltd., Bangalore, India

Research

Research Topic

Engineering sociotechnical systems (STSs) is the overarching theme of my research. An STS consists of interacting *principals* (humans and organizations) in the social tier and artificially intelligent (AI) *agents*, representing principals, in the technical tier. This conceptualization of an STS provides an infrastructure for addressing complex societal problems involving both human stakeholders and AI support.

I focus on three key abstractions necessary for engineering an STS:

- **Context** defines the micro-society an STS represents.

- **Values** represent a principal's deep-rooted motivations to act.
- **Norms** capture directed expectations between the principals.

Whereas these concepts are ingrained in human decision making, they are extremely challenging for AI agents to learn and incorporate in their decision making. My research endeavor is to enable AI agents to learn values, norms, and context so that they can augment humans in decision making.

In this endeavor, I have carved three intersecting technical research lines.

- **Natural Language Processing (NLP).** I develop NLP techniques, including Large Language Models (LLMs), to recognize principals' values and norms, often expressed in natural language. To do so, I rely on interaction data from social media, deliberation platforms, and stakeholder meetings.
- **Multiobjective Decision Making (MODM).** Most practical decision-making scenarios involve multiple stakeholders, and consequently, multiple objectives. First, I develop methods to formulate stakeholders' values and norms extracted via NLP as objectives. Second, I develop decision methods to support decision makers and multiobjective optimization methods to work interactively.
- **Hybrid Intelligence (HI).** I develop systematic methodologies for engineering HI systems, where agents augment, as opposed to replacing, human intelligence. My work in this line involves developing conceptual models and evaluation methods for HI systems, and modeling interdependencies.

These research lines are challenging for three main reasons. First, they involve concepts such as values and norms that are not well defined in the AI literature. Further, in the broader literature (e.g., philosophy) there are multiple perspectives on values and norms. I undertake the challenging but important task of synthesizing such perspectives and translating them into computational abstractions. This is an essential step in enabling AI agents to reason about values and norms. Second, learning values and norms from data is challenging because they are context-specific, subjective, and often, not explicit in data. Finally, I am interested in developing systems end-to-end. Thus, values are not an end for me but a means to support decision making. This expands the scope of my research substantially, e.g., by having to combine different forms of intelligence (human and artificial) and balancing diverse objectives. However, this is important because it demonstrates the practical and societal importance of my research.

Application Potential

My research is broadly applicable, i.e., not specific to a domain. The tools, techniques, and methods I develop apply to many complex problems where decision making involves multiple stakeholders.

In the recent years, I have focused on **climate change** (mitigation and adaptation) as an application domain for three reasons. First, the multiple facets of the climate change problem, provide me an opportunity to demonstrate how the research lines I am working on come together. That is, I am exploring this domain to extract values and norms from a variety of textual documents, facilitate deliberations among stakeholders, and apply multiobjective optimization methods. Second, climate change is one of the biggest and urgent problems our society is facing, and I am strongly motivated to work on societal problems. Finally, this is a strategic choice since I have strong collaborators in this domain. Specifically, in the Hippo Delft AI Lab, I am working with leading researchers on climate policies, and in the Red&Blue project, I am working with leading experts investigating the risks of climate change on the built environment. More recently, I am exploring **health and lifestyle** as an application domain. Specifically, in a recent NWO (KIC) project, I am investigating the role values, norms, and hybrid intelligence in enhancing mental health.

Publications Record

I strive for publishing in high-quality venues. This is reflected in the types of papers I have published.

- Number of peer-reviewed papers: 57.
  - Journal articles: 12.
  - Major conference proceedings: 28.
- Number of papers in venues rated A* (**Exceptional**): 14.
  - Conferences: AAMAS $\times$ 7, IJCAI $\times$ 2, ARR (ACL, EMLP) $\times$ 3.
  - Journals: TOCHI $\times$ 1, TOSEM $\times$ 1.
- Number of papers in venues rated A (**Excellent**): 12.
  - Conferences: RE $\times$ 5, ARR (NAACL, EACL) $\times$ 3, RecSys $\times$ 1, ESEM $\times$ 1, ICSOC $\times$ 1.

- Journals: JAAMAS $\times$ 1, JAIR $\times$ 1, Intelligent Systems $\times$ 1.

The categorization of publication venues as Exceptional (A*) or Excellent (A) is based on CORE (Computing Research and Education) portal for conference and journal rankings. In this ranking, 7.5% of ranked conferences are A* and 14.5% are A, and 6% of ranked journals are A* and 10% are A.

I have a good balance of journal and conference publications. In my field, conference publications are archival, and typically, as prestigious as (sometime more prestigious than) journal publications. Accordingly, I often publish in top-tier conference proceedings.

The topics of my research are of interest to several computing disciplines. Inline with the three research directions above, I pursue publication in two main categories of venues: (1) NLP venues such as ACL, EMNLP, NAACL, and EACL, and (2) Multiagent Systems venues such as AAMAS, IJCAI, JAAMAS, and JAIR. From time to time, I also publish in Software Engineering and HCI venues.

According to Google Scholar, the following are my key citation metrics.

- Number of citations: 1132.
- h-index: 20.
- i10 index: 35.

The topics I work on are gaining increasing traction as reflected in the upward trend in my citations.

Awards and Recognition

- **Best Paper Award** and **Best Poster Award** — June 2022
  First International Conference on Hybrid Human-Artificial Intelligence (HHAI)
  Paper Title: HyEnA: A Hybrid Method for Extracting Arguments from Opinions
  Poster Title: Identifying Context-Specific Values via Hybrid Intelligence
- **Best Blue Sky Paper Award** — May 2020
  19th Intl. Conf. on Autonomous Agents and Multi-Agent Systems (AAMAS)
  Title: New Foundations of Ethical Multiagent Systems
- **Outstanding Dissertation Award** — May 2017
  Department of Computer Science, North Carolina State University
  Title: Engineering personal agents: Toward personalized, context-aware, and privacy-preserving applications
- **Outstanding Research Award** — Apr 2016
  Department of Computer Science, North Carolina State University
  Awarded to one or two graduate students in the department every year
- **Best Program Committee Member Award** — May 2022
  21st Intl. Conf. on Autonomous Agents and Multi-Agent Systems (AAMAS)
  Awarded to 14 out of 487 PC members

Five Key Publications

The five publications below exemplify the research abstractions (values, norms, and context) and the research lines (NLP, MODM, and HI) my works focus on. These works demonstrate my **visionary** ideas, **breadth and depth** of my research, their scientific **impact**, and my research **growth**.

1. **A Hybrid Intelligence Method for Argument Mining**
   M van der Meer, E Liscio, CM Jonker, A Plaat, P Vossen, and PK Murukannaiah.
   Journal of Artificial Intelligence Research (JAIR), pages 1–33, To appear, 2024.

2. **What Lies beyond the Pareto Front?**
   A Survey on Decision-Support Methods for Multi-Objective Optimization?
   Z Osika, J ZatarainSalazar, DM Roijers, FA Oliehoek, and PK Murukannaiah.
   In Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI), pages 6741-6749, Macao S.A.R, 2023.

3. **What Values should an Agent Align with?**
   An Empirical Comparison of General and Context-Specific Values.
   E Liscio, M van der Meer and LC Siebert and CM Jonker and PK Murukannaiah.
   Autonomous Agents and Multi-Agent Systems (JAAMAS), 36(1):1–23, 2022.

4. **New Foundations of Ethical Multiagent Systems.**
   PK Murukannaiah, N Ajmeri, CM Jonker, and MP Singh.
   In Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS), pages 1706–1710, Auckland, 2020.

5. **Xipho: Extending Tropos to Engineer Context-Aware Personal Agents.**
   PK Murukannaiah and MP Singh.
   In Proceedings of the 13th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS), pages 309–316, Paris, 2014.

Visionary ideas: Murukannaiah et al. [4] is an AAMAS Blue Sky paper that outlines my vision for engineering sociotechnical systems with values, norms, and context at the core of my vision. The AAMAS Blue Sky track selects papers that "provoke visionary ideas, long-term challenges, new research opportunities, and controversial debate." Osika et al. [2] is an IJCAI Survey track paper that outlines my vision on multi-objective decision making on societal issues such as climate change. The IJCAI Survey track selects papers that "describe how lessons learned from the topic can contribute to new ideas and visions that can stimulate the research community to pursue new directions."

Breadth and depth: The selected papers demonstrate the breadth of my research across the three research lines, e.g., van der Meer et al. [1] is good example for my work on NLP, Liscio et al. [3] for HI, and Osika et al. [2] for MODM. Further, these works demonstrate how my research lines intersect, e.g., van der Meer et al. [1] includes aspects of both NLP and HI, and Osika et al. [2] includes aspects of both MODM and HI. Finally, these works also demonstrate the depth of my work. These are not one-off papers, but have been preceded or followed by additional work. I have published three more papers with Liscio, three more papers with van der Meer, and one more paper with Osika based on these topics.

Scientific impact: My works have been published in highly selective venues. The five papers in this list are published in venues that are rated exceptional (A*) or excellent (A) in my field. Further, Murukannaiah et al. [4] received the Best Blue Sky Paper award at AAMAS 2020 and the conference version of van der Meer et al. [1] received the Best Paper award at HHAI 2022. This shows that my works have been recognized as high-quality by the scientific community. In addition, my works are also being increasingly cited. For example, Murukannaiah et al. [4] has received 60 citations within four years, and Liscio et al. [3] and its conference version, together, have received 71 citations within three years.

Research growth: Murukannaiah and Singh [5] is my first major conference paper and it laid foundation for my PhD. Contrasting [5] with the other papers in the list shows, simultaneously, that the core aspects of my thinking have remained the same, but several new research directions emerged, over the years.

Five Key Artifacts

In addition to publishing scientific papers, my research output also includes software and datasets, which assist in broad dissemination of my work. Below, I include five such key artifacts and for each, I refer to a paper that provides additional details on the artifact.

- **HyEnA: A Hybrid Intelligence System for Argument Mining.** Code and dataset
  - M. van der Meer, E. Liscio, C. M. Jonker, A. Plaat, P. Vossen, and P. K. Murukannaiah. HyEnA: A hybrid method for extracting arguments from opinions. In *Proceedings of the First International Conference on Hybrid Human-Artificial Intelligence*, HHAI, pages 17–31, Amsterdam, 2022

- **Axies: A Collaborative Platform for Identifying Context-Specific Values.** Software and Demo
  - E. Liscio, M. van der Meer, C. M. Jonker, and P. K. Murukannaiah. A collaborative platform for identifying context-specific values. In *Proceedings of the 20th Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '21, pages 1773–1775, London, 2021a

- **Crowd RE: A Dataset for Crowd-Based Requirements Engineering.** Dataset
  - P. K. Murukannaiah, N. Ajmeri, and M. P. Singh. Toward automating Crowd RE. In *IEEE 25th International Requirements Engineering Conference*, pages 512–515, Lisbon, Sept. 2017a

- **PrIncipedia: A Privacy Incidents Encyclopedia.** Database snapshot
  - P. K. Murukannaiah, J. Staddon, H. Lipford, and B. Knijnenburg. PrIncipedia: A privacy incidents

encyclopedia. In *The 9th Annual Privacy Law Scholars Conference*, 2016c

- **Platys: A Middleware for Developing Context-aware Personal Agents.** Software
  - P. K. Murukannaiah, R. Fogues, and M. P. Singh. Platys: A framework for supporting context-aware personal agents. In *Proceedings of the 13th International Conference on Autonomous Agents and Multi-agent Systems*, pages 1689–1690, Paris, 2014

Keynote Talks

- **Values and Norms as the Ethical Foundations of Sociotechnical Systems** Dec 2022
  International Conference on Data Science, Agents and AI (*Chennai, India / Online*)
- **Enabling Privacy Analytics via a Privacy Incidents Database** Aug 2017
  International Conference on Privacy, Security and Trust (*Calgary, Canada*)

Invited Talks in Seminars

- **On the Value Alignment of Large Language Models** Apr 2024
  Workshop on Large Language Models for media and democracy: wrecking or saving society? (*Amsterdam*)
- **From Deliberations to Decisions via Hybrid Intelligence** Apr 2024
  Citizen-Centric AI Seminar Series (*Raleigh, USA / Online*)
- **Estimating Context Specific Values from Natural Language** Oct 2022
  AI in Society Seminar Series (*Online*)
- **Multiagent Foundations for AI Ethics** May 2021
  AI Tech Agora (*Delft*)
- **Engineering Intelligent Agents** Sep 2017
  Rochester Institute of Technology (*Rochester, USA*)
- **RE in the Age of Mobile and Social** Mar 2017
  University of Ottawa (*Ottawa, Canada*)

Advisees

Postdoctoral Research Associates (Current)

1. **Enrico Liscio** Apr 2024–Mar 2027
   Intelligent Systems, EEMCS, TU Delft
   Topic: Analyzing Public Values in Media via NLP
2. **Edgar Salas Gironés** Sep 2023–Aug 2026
   Intelligent Systems, EEMCS, TU Delft
   Topic: Artificial Intelligence for Collaborative Climate Risk Management

Doctoral Students (Current)

1. **Michiel van der Meer** Jun 2020–Jul 2024
   Leiden Institute of Advanced Computer Science (LIACS), Leiden University
   Working Title: Facilitating Online Opinion Diversity through Hybrid NLP Approaches
   Role: Daily Supervisor
2. **Zuzanna Osika** Apr 2022–Mar 2027
   Intelligent Systems, EEMCS, TU Delft
   Topic: Explainable Multi-Objective Decision Support Methods
   Role: Daily Supervisor
3. **Shubhalaxmi Mukherjee** Aug 2023–Jul 2028
   Intelligent Systems, EEMCS, TU Delft
   Topic: Fact Checking using Large Language Models
   Role: Daily Supervisor

4. **Feline L. Lindeboom** Sep 2023–Oct 2027
   Bernoulli Institute for Maths, CS and AI; University of Groningen
   Topic: Algorithm Supported Deliberation
   Role: Co-Supervisor (with Dr. Davide Grossi)
5. **Anastasia Akkuzu** Oct 2023–Sep 2027
   Information and Computing Sciences; Utrecht University
   Topic: Responsible Autonomy for Hybrid Intelligence
   Role: Co-Supervisor (with Prof. Dr. Pinar Yolum)

Doctoral Students (Graduated)

1. **Enrico Liscio** Mar 2020–Apr 2024
   PhD in Computer Science (**Cum Laude**)
   Context-Specific Value Inference via Hybrid Intelligence
   Role: Daily Supervisor

MSc Thesis Students (Current)

1. **Joost Jansen** Nov 2023–Jul 2024
   MSc in Computer Science
   Topic: Using NLP to Support Responsible Investing
2. **Budi Han** Jan 2024–Sep 2024
   MSc in Computer Science
   Topic: AI Health Assistant

MSc Thesis Students (Graduated)

1. **Jahson O'Dwyer Wha Binda** Oct 2023–May 2024
   MSc in Computer Science
   Transformer Modules
2. **Jakub Nguyen** Nov 2022–Aug 2023
   MSc in Computer Science, TU Delft
   The generalizability of argument quality dimensions in NLP models
3. **Jeongwoo Park** Nov 2022–Jul 2023
   MSc in Computer Science, TU Delft
   Constructing a Pluralist Moral Sentence Embedding Space using Contrastive Learning
4. **Mei Lan Schrama** Dec 2021–Sep 2022
   MSc in Computer Science, TU Delft
   Active Learning with Multi-annotator Disagreement
5. **Mila Hendrikse** Feb 2021–Nov 2021
   MSc in Industrial Ecology, TU Delft (Co-Supervisor)
   Unravelling Twitter chaos during a policy crisis
6. **Divyae Mittal** Feb 2020–Jul 2020
   MSc in Architecture, Urbanism and Building Sciences, TU Delft (External Supervisor)
   Simulation of adaptive user behavior with respect to noises in offices
7. **Arvind Sastha Kumar** Jan 2020–Aug 2021
   MSc in Computer Science, NC State University (Co-Supervisor)
   Federated Inverse Averaging: Toward Pareto-Optimal Minima in Human Activity Recognition
8. **Bhavana Balraj** Jan 2020–Aug 2021
   MSc in Computer Science, NC State University (Co-Supervisor)
   Multilabel Active Learning for User Context Recognition in the Wild
9. **Aditya Parkhi** Jan 2020–Aug 2021
   MSc in Computer Science, NC State University (Co-Supervisor)
   Novelty Detection in Textual Data using Document Encoders and Self Attention
10. **Venkatesh Thimma Dhinakaran** Jan 2019–Sep 2019
    MSc in Computer Science, Rochester Institute of Technology
    App Review Analysis via Active Learning:
    Reducing Supervision Effort Without Compromising Classification Accuracy

Teaching

Delft University of Technology

- **Advanced Interdisciplinary Artificial Intelligence Project (AI2P) – 15 ECTS**
  MSc course available to students from all faculties
  - Course Developer — 2021
  - Course Coordinator — Q1 2022, 2023
- **NLP for Society (Under development) – 5 ECTS**
  MSc in DSAIT (Theme: Language and Speech Technology)
  - Course Developer — 2024
  - Course Coordinator and Lecturer — Q4 2025
- **Collaborative Artificial Intelligence (CAI) – 5 ECTS**
  BSc in Computer Science (Third year)
  - Course Developer — 2020
  - Course Coordinator and Lecturer — Q3 2021–2024
- **Computational Intelligence (CI) – 5 ECTS**
  BSc in Computer Science (Second year)
  - Course Coordinator — Q3 2020, 2021
  - Lecturer — Q3 2020–2024
- **Conversational Agents – 5 ECTS**
  MSc in Computer Science
  - Lecturer — Q3 2021
- **Research Methods in Data Science – 5 ECTS**
  MSc in Computer Science
  - Lecturer — Q3 2020

Rochester Institute of Technology

- **Data Science Methods for Software Engineering**
  MS in Software Engineering
  - Course Developer — 2017
  - Course Coordinator and Lecturer — Spring 2017, 2018
- **Engineering Secure Software**
  BS in Software Engineering
  - Course Coordinator and Lecturer — Spring 2017, 2018

Technical Tutorials

- **Ethics in Sociotechnical Systems** — 2020–2022
  Delivered at AAAI, IJCAI, and AAMAS

Valorization

Funded Projects (TU Delft)

I was involved in writing the proposal for each of the following projects. The budget (or number pf PhDs or PostDocs) mentioned is the share for the Intelligent Systems department of TU Delft.

- **Technology Assisted Self-Management: Preventing Relapse and Crisis by the Severe Mentally Ill Themselves** — €354,928 — Set 2023–Dec 2029
  Knowledge and Innovation Covenant (NWO-KIC)
  Role: I was a co-applicant of the proposal. I co-supervisor a Post-Doc and a Research Engineer in the project.
- **RED&BLEUE: Real Estate Development & Building in Low Urban Environments** — €228,698 — Sep 2022–Dec 2027
  Dutch Research Agenda (NWA)
  Role: I was a co-applicant of the proposal. I co-lead a work package and co-supervisor a PostDoc in the project.

- **Hippo Delft AI Lab:** 2 PhDs Sep 2021–Dec 2026
  AI for Fair, Efficient, and Interpretable Policy Analysis
  TU Delft AI Initiative
  Role: I am one of the two Co-PIs of the proposal and one of the two co-directors of the Hippo lab. I supervise two PhDs in the lab.
- **Urban Dialogues:** €45,000 Sep 2022–Dec 2024
  Dialogue-Driven Co-Creation to Support Citizen Science
  BOLD Cities Network
  Role: I am one of the four Co-PIs of the proposal. I supervise one MSc student in the project.

Contributing Projects (TU Delft)

I joined the following projects after they were funded (I was not involved in writing the proposal). The number of PhDs or PostDocs mentioned is the share for the Intelligent Systems department of TU Delft.

- **Hybrid Intelligence: Augmenting Human Intellect** 3 PhDs Jan 2020–Dec 2029
  NWO Zwaartekracht
  Role: I have important leadership roles in the consortium, e.g., research line coordinator (past) and use cases coordinator (present). I supervise three PhDs (one as daily supervisor and two as co-supervisor). I represented the project in mid-term renewal.
- **ALGOSOC: Public Values in the Algorithmic Society** 1 PostDoc Oct 2022–Oct 2032
  NWO Zwaartekracht
  Role: I supervise a PostDoc in the project.

Funded Projects (NC State University)

I was involved in writing the proposal for the following project as an external collaborator, i.e., there was no direct budget allocated to me, but I co-supervised students in the project.

- **Principles of Normative Multiagent Systems for Decentralized Applications** – Aug 2019–May 2023
  US National Science Foundation
  Role: I lead one of the tasks in the project and co-supervise a PhD.

Funded Projects (Rochester Institute of Technology)

The following projects, where I am a PI/Co-PI/Senior Personnel, were funded when I was an Assistant Professor at RIT. When I left RIT, I transferred these projects to my colleagues, and collaborated externally.

- **REU Site: Cultivating Next Generation Software Engineering Researchers** $360,000 Mar 2018–Feb 2021
  US National Science Foundation
  Role: Senior Personnel
- **Citizenly: Empowering Communities by Democratizing Urban Data Science** $298,734 Oct 2019–Sep 2021
  US National Science Foundation
  Role: Co-PI (Did not participate since I left RIT)
- **Facilitating Creativity in Crowd RE** $14,850 Mar 2017–Jan 2018
  Rochester Institute of Technology
  Role: PI
- **PrIncipedia: A Privacy Incidents Community Platform** $5,000 May 2019–Aug 2020
  Rochester Institute of Technology
  Role: PI

Leadership

Hybrid Intelligence (HI) Consortium

- **Use Cases Coordinator** 2023–Present
  The four HI use cases (Robotic Surgery, Diabetes Management, Scientific Assistant, and Education) showcase hybrid intelligence in practical applications. Each case study has leaders, and I oversee all four case studies. I assist case study leaders in developing the vision, foster connections across case studies, and connect the case studies to the external world.
- **Diversity Committee Co-Chair** 2023–Present
  The motto of the HI Diversity committee is to create "Hybrid Intelligence for All." As part of the three-member HI diversity committee, I develop guidelines for inclusive recruitment, assess the diversity of the consortium, and influence the recruitment process.
- **Research Line Coordinator** 2020–2022
  The HI consortium was initially organized along four research lines: Collaborative, Adaptive, Responsible, Explainable. I was one of the two coordinators of the Explainable HI research line. In this role, I developed a research agenda for the research line and evaluated the progress of multiple projects in this research line along the HI dimensions.

Research Community

- **Program Co-Chair** 2024
  3rd International Conference on Hybrid Human-Artificial Intelligence (HHAI)
- **Co-Editor** 2020–Present
  IEEE Internet Computing, Internet Ethics Department
- **Competitions and Challenges Track Co-Chair** 2023, 2024
  32nd International Joint Conference on Artificial Intelligence (IJCAI)
- **Workshops and Tutorials Track Co-Chair** 2023
  2nd International Conference on Hybrid Human-Artificial Intelligence (HHAI)
- **Student Volunteer Co-Chair** 2018
  26th IEEE International Requirements Engineering Conference, August 20–24, 2018
- **Workshop Co-Chair**
  - Fifth Intl. Workshop on Artificial Intelligence for Requirements Engineering, Banff 2018
  - Second International Workshop on Crowd-Based Requirements Engineering, Lisbon 2017
  - First International Workshop on Social Media Analytics for Smart Cities, Singapore 2017
- **PhD Symposium Co-Chair** 2016
  Doctoral Symposium on Foundations and Applications of Self-∗ Systems (FAS∗), Augsburg, Germany, September 12-16, 2016
- **Senior Program Committee Member**
  - International Joint Conference on Artificial Intelligence (IJCAI) 2023, 2024
  - International Conf. on Autonomous Agents and MultiAgent Systems (AAMAS) 2023, 2024
- **Program Committee (PC) Membership (Selected)**
  - European Conference on Artificial Intelligence (ECAI) 2023, 2024
  - International Conference on Hybrid Human-Artificial Intelligence (HHAI) 2022, 2023
  - International Joint Conference on Artificial Intelligence (IJCAI) 2023, 2024
  - International Conf. on Autonomous Agents and MultiAgent Systems (AAMAS) 2018–2022
  - ACM Conference on User Modeling, Adaptation and Personalization (UMAP) 2022, 2023
  - AAAI Conference on Artificial Intelligence (AAAI) 2021, 2022
- **Journal Information Director** 2013–2018
  ACM Transactions on Internet Technology

Organization

- **Program Coordinator** 2024–Present
  MSc in Data Science and AI Technology (DSAIT)
  - I am one of the two Master Coordinators for the MSc DSAIT program. In this role, under the overall leadership of the Director of Studies, I am responsible for (1) coordinating the implementation of

education policy and the program curriculum, and quality assurance, (2) advising students and evaluating their IEPs, (3) monitoring the progress of the students, and (4) reporting progress to Board of Education and Management team. As I am looking forward to the program kickoff, I am currently mastering all relevant policies, procedures, and tools.

- **Education Liason, TU Delft AI L&T Program** 2021–Present
  TU Delft AI Labs & Talent Program Management Team
  - I am one of the five members of the TU Delft AI Labs & Talent Program Management Team. The L&T M-Team, under the overall leadership of the Pro Vice Rector Magnificus, oversees the 24 Delft AI labs and additional talents, committed to education, research, and innovation in AI, Data and Digitalisation. In this team, I serve as the Education Liaison. My key responsibilities include (1) facilitating a steady dialogue between L&T members and AI Education teams, to identify needs, opportunities or priorities for AI education, (2) contributing to AI Education with wishes/ideas coming from L&T, and (3) linking L&T activities to relevant local faculty activities.
- **MSc DSAIT Program Development Team** 2021–2023
  MSc in Data Science and AI Technology (*DSAIT*)
  - I was a member of the core team (16 members) that lead the development of MSc in DSAIT. As part of this team, I (1) investigated similar programs in other universities and reference frameworks, e.g., from KION and ACM, for positioning DSAIT, (2) developed the intended learning outcomes of the program, (3) developed a program structure based on the reference frameworks and the available teaching resources, and (4) represented the program to the NVAO accreditation panel. The program has been accredited and I have been entrusted with the role of Master Coordinator.
- **Interdisciplinary AI Education Exploration Team** 2020–2021
  FAIP and AI2P MSc AI Blocks
  - In a three-member working group, I explored the options for creating university-wide training programs that addresses common AI needs across MSc programs in all faculties. Our working group was conceived with a vision that societal applications of AI require both fundamental AI skills as well as domain experts that can work with AI experts. Our working group proposed two 15 EC MSc AI Blocks: Foundations of AI Program (FAIP) and Advanced Interdisciplinary AI Project (AI2P), considering the breadth of AI pre-requisites AI knowledge students from different faculties may have. These programs were further developed by the course coordinators (I coordinate AI2P). Both programs are running for two years now, attracting students from multiple faculties.

All Publications

Dissertation

1. P. K. Murukannaiah. *Engineering personal agents: Toward personalized, context-aware, and privacy-preserving applications*. PhD thesis, North Carolina State University, Raleigh, NC, 2016

Journal Articles (Peer-Reviewed)

2. M. van der Meer, E. Liscio, C. M. Jonker, A. Plaat, P. Vossen, and P. K. Murukannaiah. A hybrid intelligence method for argument mining. *Journal of Artificial Intelligence Research*, 80:1–36, jul 2024a. ISSN 1076-9757. doi: 10.1613/jair.1.15135

3. R. X. Lera-Leri, E. Liscio, F. Bistaffa, C. M. Jonker, M. Lopez-Sanchez, P. K. Murukannaiah, J. A. Rodriguez-Aguilar, and F. Salas-Molina. Aggregating value systems for decision support. *Knowledge-Based Systems*, 287:111453, 2024. ISSN 0950-7051. doi: https://doi.org/10.1016/j.knosys.2024.111453. URL https://www.sciencedirect.com/science/article/pii/S0950705124000881

4. R. Shortall, A. Itten, M. v. d. Meer, P. Murukannaiah, and C. Jonker. Reason against the machine? future directions for mass online deliberation. *Frontiers in Political Science*, 4, 2022. ISSN 2673-3145. doi: 10.3389/fpos.2022.946589. URL https://www.frontiersin.org/articles/10.3389/fpos.2022.946589

5. G. Yuan, M. P. Singh, and P. K. Murukannaiah. An interpretable framework for investigating the neighborhood effect in poi recommendation. *PLOS ONE*, 16(8):1–19, 08 2021. doi: 10.1371/journal.pone.0255685. URL https://doi.org/10.1371/journal.pone.0255685

6. E. Liscio, M. van der Meer, L. C. Siebert, C. M. Jonker, and P. K. Murukannaiah. What values should an agent align with? an empirical comparison of general and context-specific values. *Autonomous Agents and Multi-Agent Systems*, 36(1):23, 2022b. doi: 10.1007/s10458-022-09550-0

7. I. Kola, P. K. Murukannaiah, C. M. Jonker, and B. van Riemsdijk. Towards social situation awareness in support agents. *IEEE Intelligent Systems*, 37(5):50–58, 2022. doi: 10.1109/MIS.2022.3163625

8. H. Sapkota, P. K. Murukannaiah, and Y. Wang. A network-centric approach for estimating trust between open source software developers. *PLOS ONE*, 14(12):1–30, 12 2019. doi: 10.1371/journal.pone.0226281. URL https://doi.org/10.1371/journal.pone.0226281

9. R. L. Fogues, P. K. Murukannaiah, J. M. Such, and M. P. Singh. SoSharP: Recommending sharing policies in multiuser privacy scenarios. *IEEE Internet Computing*, 21(6):28–36, 2017a

10. R. L. Fogues, P. K. Murukannaiah, J. M. Such, and M. P. Singh. Sharing policies in multiuser privacy scenarios: Incorporating context, preferences, and arguments in decision making. *ACM Transactions on Computer-Human Interaction*, 24(1):1–29, Mar. 2017b. ISSN 1073-0516. doi: 10.1145/3038920. URL http://doi.acm.org/10.1145/3038920

11. P. K. Murukannaiah and M. P. Singh. Platys: An active learning framework for place-aware application development and its evaluation. *ACM Transactions on Software Engineering and Methodology*, 24(3): 1–33, May 2015

12. L. Zavala, P. K. Murukannaiah, N. Poosamani, T. Finin, A. Joshi, I. Rhee, and M. P. Singh. Platys: From position to place-oriented mobile computing. *AI Magazine*, 36(2):50–62, July 2015

13. P. K. Murukannaiah and M. P. Singh. Platys Social: Relating shared places and private social circles. *IEEE Internet Computing*, 16(3):53–59, May 2012. ISSN 1089-7801. doi: http://doi.ieeecomputersociety.org/10.1109/MIC.2011.106

Major Conference Papers (Peer-Reviewed)

These articles are rigorously peer-reviewed, archival, and typically published in highly selective venues.

14. D. Dell'Anna, P. K. Murukannaiah, B. Dudzik, D. Grossi, C. M. Jonker, C. Oertel, and P. Yolum. Toward a quality model for hybrid intelligence teams. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*, pages 1–10, Auckland, 2024. To appear

15. M. van der Meer, P. Vossen, C. M. Jonker, and P. K. Murukannaiah. An empirical analysis of diversity in argument summarization. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2028–2045, Malta, 2024b

16. J. Park, E. Liscio, and P. K. Murukannaiah. Morality is non-binary: Building a pluralist moral sentence embedding space using contrastive learning. In *Findings of the 18th Conference of the European Chapter of the Association for Computational Linguistics*, pages 654–673, Malta, 2024

17. M. van der Meer, P. Vossen, C. M. Jonker, and P. K. Murukannaiah. Do differences in values influence disagreements in online discussions? In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 15986–16008, Singapore, 2023

18. E. Liscio, O. Araque, L. Gatti, I. Constantinescu, C. M. Jonker, K. Kalimeri, and P. K. Murukannaiah. What does a text classifier learn about morality? An explainable method for cross-domain comparison of moral rhetoric. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics*, pages 14113–14132, Toronto, 2023a

19. Z. Osika, J. ZatarainSalazar, D. M. Roijers, F. A. Oliehoek, and P. K. Murukannaiah. What lies beyond the Pareto front? A survey on decision-support methods for multi-objective optimization. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence*, IJCAI '23, pages 6741–6749, Macao, S.A.R, 2023

20. E. Liscio, R. Lera-Leri, F. Bistaffa, R. I. J. Dobbe, C. M. Jonker, and P. K. Murukannaiah. Value inference in sociotechnical systems. In *Proceedings of the 22nd International Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '23, pages 1774–1780, London, 2023b

21. P. K. Murukannaiah, N. Ajmeri, and M. P. Singh. Enhancing creativity as innovation via asynchronous crowdwork. In *Proceedings of the 14th ACM Web Science Conference*, WebSci '22, pages 66–74, Barcelona, 2022

22. E. Liscio, A. E. Dondera, A. Geadău, C. M. Jonker, and P. K. Murukannaiah. Cross-domain classification of moral values. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 2727–2745, Seattle, 2022a

23. M. van der Meer, E. Liscio, C. M. Jonker, A. Plaat, P. Vossen, and P. K. Murukannaiah. HyEnA: A hybrid method for extracting arguments from opinions. In *Proceedings of the First International Conference on Hybrid Human-Artificial Intelligence*, HHAI, pages 17–31, Amsterdam, 2022

24. L. C. Siebert, E. Liscio, P. K. Murukannaiah, L. Kaptein, S. Spruit, J. van den Hoven, and C. M. Jonker. Estimating value preferences in a hybrid participatory system. In *Proceedings of the First International Conference on Hybrid Human-Artificial Intelligence*, HHAI, pages 114–127, Amsterdam, 2022

25. M. Tejedor-Romero, P. K. Murukannaiah, J. M. Giménez-Guzmán, I. Marsá-Maestre, and C. M. Jonker. Comparing mediated and unmediated agent-based negotiation in wi-fi channel assignment. In *Proceedings of 24th International Conference on Principles and Practice of Multi-Agent Systems (PRIMA)*, volume 13753 of *Lecture Notes in Computer Science*, pages 592–601. Springer, 2022b. doi: 10.1007/978-3-031-21203-1_37. URL https://doi.org/10.1007/978-3-031-21203-1_37

26. E. Liscio, M. van der Meer, L. C. Siebert, C. M. Jonker, N. Mouter, and P. K. Murukannaiah. Axies: Identifying and evaluating context-specific values. In *Proceedings of the 20th Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '21, pages 799–808, London, 2021b

27. S. Bennati, C. M. Jonker, P. K. Murukannaiah, R. Shinde, and T. Verwaart. Discrimination between social groups: The influence of inclusiveness-enhancing mechanisms on trade. In *Proceedings of the 11th International Conference on Simulation and Modeling Methodologies, Technologies and Applications*, SimulTech, pages 1–12, 2021

28. D. Mittal, M. Turrin, M. J. Tenpierik, P. K. Murukannaiah, and R. Goldstein. A framework for multi-agent and acoustic simulation for office design. In *To apperar in the Proceedings of the 12th Annual Symposium on Simulation for Architecture and Urban Design*, SimAUD '21, pages 1–9, 2021

29. P. K. Murukannaiah, N. Ajmeri, C. M. Jonker, and M. P. Singh. New foundations of ethical multiagent systems. In *Proceedings of the 19th Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '20, pages 1706–1710, Auckland, 2020a

30. N. Ajmeri, H. Guo, P. K. Murukannaiah, and M. P. Singh. Elessar: Ethics in norm-aware agents. In *Proceedings of the 19th Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '20, pages 16–24, Auckland, 2020

31. M. Alshangiti, H. Sapkota, P. K. Murukannaiah, X. Liu, and Q. Yu. Why is developing machine learning applications challenging? a study on stack overflow posts. In *Proceedings of the ACM/IEEE International Symposium on Empirical Software Engineering and Measurement*, ESEM '19, pages 1–11, Porto de Galinhas, Brazil, Sept. 2019

32. V. T. Dhinakaran, R. Pulle, N. Ajmeri, and P. K. Murukannaiah. App review analysis via active learning: Reducing supervision effort without compromising classification accuracy. In *Proceedings of the 26th IEEE International Requirements Engineering Conference*, RE '18, pages 170–181, Banff, Canada, Aug. 2018

33. N. Ajmeri, H. Guo, P. K. Murukannaiah, and M. P. Singh. Robust norm emergence by revealing and reasoning about context: Socially intelligent agents for enhancing privacy. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, IJCAI '18, pages 28–34, Stockholm, July 2018b

34. N. Ajmeri, P. K. Murukannaiah, H. Guo, and M. P. Singh. Arnor: Modeling social intelligence via norms to engineer privacy-aware personal agents. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '17, pages 230–238, São Paulo, Brazil, 2017. URL http://dl.acm.org/citation.cfm?id=3091125.3091163

35. G. M. Kanchev, P. K. Murukannaiah, A. K. Chopra, and P. Sawyer. Canary: Extracting requirements-related information from online discussions. In *Proceedings of the 25th IEEE International Requirements*

Engineering Conference, pages 31–40, Lisbon, Sept. 2017a

36. P. K. Murukannaiah, N. Ajmeri, and M. P. Singh. Acquiring creative requirements from the crowd: Understanding the influences of personality and creative potential in Crowd RE. In *Proceedings of the 24th IEEE International Requirements Engineering Conference*, pages 176–185, Beijing, Sept. 2016b

37. G. Yuan, P. K. Murukannaiah, and M. P. Singh. Percimo: A personalized community model for location estimation in social media. In *Proceedings of the IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, pages 271–278, San Francisco, 2016

38. P. K. Murukannaiah, A. K. Kalia, P. R. Telang, and M. P. Singh. Resolving goal conflicts via argumentation-based analysis of competing hypotheses. In *Proceedings of the 23rd IEEE International Requirements Engineering Conference*, pages 156–165, Ottawa, Aug. 2015

39. A. K. Kalia, P. K. Murukannaiah, and M. P. Singh. TRACE: A dynamic model of trust for people-driven service engagements. In *Proceedings of the 13th International Conference on Service Oriented Computing*, pages 353–361, Goa, Nov. 2015

40. P. K. Murukannaiah and M. P. Singh. Xipho: Extending Tropos to engineer context-aware personal agents. In *Proceedings of the 13th International Conference on Autonomous Agents and Multi-Agent Systems*, pages 309–316, Paris, 2014a

41. G. Yuan, P. K. Murukannaiah, Z. Zhang, and M. P. Singh. Exploiting sentiment homophily for link prediction. In *Proceedings of the 8th ACM Conference on Recommender Systems*, pages 17–24, Foster City, CA, 2014. ISBN 978-1-4503-2668-1. doi: 10.1145/2645710.2645734. URL http://doi.acm.org/10.1145/2645710.2645734

Demonstration and Dataset Papers (Peer-Reviewed)

These articles are peer-reviewed and archival, but publication is not as selective as the main track.

42. E. Liscio, M. van der Meer, C. M. Jonker, and P. K. Murukannaiah. A collaborative platform for identifying context-specific values. In *Proceedings of the 20th Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '21, pages 1773–1775, London, 2021a

43. B. S. Meyers, N. Munaiah, E. Prud'hommeaux, A. Meneely, C. O. Alm, J. Wolff, and P. K. Murukannaiah. A dataset for identifying actionable feedback in collaborative software development. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, ACL '18, pages 126–131, Melbourne, July 2018

44. G. M. Kanchev, P. K. Murukannaiah, A. K. Chopra, and P. Sawyer. Canary: An interactive and query-based approach to extract requirements from online forums. In *Proceedings of the 25th IEEE International Requirements Engineering Conference*, pages 470–471, Sept. 2017b. doi: 10.1109/RE.2017.84

45. P. K. Murukannaiah, N. Ajmeri, and M. P. Singh. Toward automating Crowd RE. In *IEEE 25th International Requirements Engineering Conference*, pages 512–515, Lisbon, Sept. 2017a

46. N. Munaiah, A. Meneely, and P. K. Murukannaiah. A domain-independent model for identifying security requirements. In *IEEE 25th International Requirements Engineering Conference*, pages 506–511, Sept. 2017a. doi: 10.1109/RE.2017.79

47. P. K. Murukannaiah, R. Fogues, and M. P. Singh. Platys: A framework for supporting context-aware personal agents. In *Proceedings of the 13th International Conference on Autonomous Agents and Multi-agent Systems*, pages 1689–1690, Paris, 2014

Workshop and Other Conference Papers (Peer-Reviewed)

These articles are peer-reviewed and archival, but publication is not as selective as the main track.

48. A. Agiollo, L. C. Siebert, P. K. Murukannaiah, and A. Omicini. The quarrel of local post-hoc explainers for moral values classification in natural language processing. In *Proceedings of the 5th International Workshop on Explainable and Transparent AI and Multi-Agent Systems (EXTRAAMAS)*, volume 14127 of *Lecture Notes in Computer Science*, pages 97–115, London, 2023. Springer. doi: 10.1007/978-3-031-40878-6_6. URL https://doi.org/10.1007/978-3-031-40878-6_6

49. M. Tejedor-Romero, P. K. Murukannaiah, J. M. Giménez-Guzmán, I. Marsá-Maestre, and C. M. Jonker. Distributed multi-agent negotiation for wi-fi channel assignment. In *Proceedings of Recent Advances in Agent-Based Negotiation: Applications and Competition Challenges (ACAN)*, volume 1092 of *Studies in Computational Intelligence*, pages 3–14. Springer, 2022a. doi: 10.1007/978-981-99-0561-4_1. URL https://doi.org/10.1007/978-981-99-0561-4_1

50. M. F. Rahman, N. Sharma, and P. K. Murukannaiah. Predicting the infuence of urban vacant lots on neighborhood property values. In *Proceedings of the 1st International Conference on Urban Data Science*, UDS '20, pages 1–16, Chennai, 2020

51. G. M. Kanchev, P. K. Murukannaiah, and A. K. Chopra. Crowd-informed goal models. In *Proceedings of the Fifth International Workshop on Artificial Intelligence for Requirements Engineering*, AIRE '18, pages 47–53, Banff, Canada, Aug. 2018

52. P. K. Murukannaiah, C. Dabral, K. Sheshadri, E. Sharma, and J. Staddon. Learning a privacy incidents database. In *Proceedings of the Hot Topics in Science of Security: Symposium and Bootcamp*, HoTSoS, pages 35–44, 2017b. ISBN 978-1-4503-5274-1. doi: 10.1145/3055305.3055309. URL http://doi.acm.org/10.1145/3055305.3055309

53. P. K. Murukannaiah, J. Staddon, H. Lipford, and B. Knijnenburg. Is this a privacy incident? Using news exemplars to study end user perceptions of privacy incidents (work in progress). In *Proceedings of the Workshop on Usable Security*, USEC, pages 1–7, 2017c

54. N. Munaiah, B. S. Meyers, C. O. Alm, A. Meneely, P. K. Murukannaiah, E. Prud'hommeaux, J. Wolff, and Y. Yu. Natural language insights from code reviews that missed a vulnerability - A large scale study of chromium. In *Proceedings of the 9th International Symposium on Engineering Secure Software and Systems*, pages 70–86, 2017b. doi: 10.1007/978-3-319-62105-0_5. URL https://doi.org/10.1007/978-3-319-62105-0_5

55. P. K. Murukannaiah, J. Staddon, H. Lipford, and B. Knijnenburg. PrIncipedia: A privacy incidents encyclopedia. In *The 9th Annual Privacy Law Scholars Conference*, 2016c

56. R. L. Fogues, P. K. Murukannaiah, J. M. Such, A. Espinosa, A. Garcia-Fornes, and M. P. Singh. Argumentation for multi-party privacy management. In *Second International Workshop on Agents and CyberSecurity*, pages 3–6, 2015

57. P. K. Murukannaiah. Reasoning about context and engineering context-aware agents. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems*, AAMAS '14, pages 1733–1734, 2014. ISBN 978-1-4503-2738-1. URL http://dl.acm.org/citation.cfm?id=2615731.2616150

58. C.-W. Hang, P. K. Murukannaiah, and M. P. Singh. Platys: User-centric place recognition. In *AAAI Workshop on Activity Context-Aware Systems*, 2013

Technical Columns

59. M. P. Singh and P. K. Murukannaiah. Toward an ethical framework for smart cities and the internet of things. *IEEE Internet Computing*, 27(2):51–56, 2023. doi: 10.1109/MIC.2023.3236826

60. I. Marsa-Maestre, J. M. Gimenez-Guzman, M. Tejedor-Romero, E. de la Hoz, and P. K. Murukannaiah. Democratic wireless channel assignment: Fair resource allocation in wi-fi networks. *IEEE Internet Computing*, 27(01):76–80, jan 2023. ISSN 1941-0131. doi: 10.1109/MIC.2022.3201454

61. P. K. Murukannaiah and M. P. Singh. From machine ethics to internet ethics: Broadening the horizon. *IEEE Internet Computing*, 24(3):51–57, 2020

62. N. Ajmeri, H. Guo, P. K. Murukannaiah, and M. P. Singh. Designing ethical personal agents. *IEEE Internet Computing*, 22(2):16–22, Mar. 2018a. ISSN 1089-7801. doi: 10.1109/MIC.2018.022021658

63. P. K. Murukannaiah, N. Ajmeri, and M. P. Singh. Engineering privacy in social applications. *IEEE Internet Computing*, 20(2):72–76, Mar. 2016a

64. P. K. Murukannaiah and M. P. Singh. Understanding location-based user experience. *IEEE Internet Computing*, 18(6):72–76, Nov. 2014b. ISSN 1089-7801. doi: 10.1109/MIC.2014.127

Edited Proceedings

65. P. K. Murukannaiah and T. Hirzle, editors. *Proceedings of the Workshops at the Second International Conference on Hybrid Human-Artificial Intelligence co-located with (HHAI 2023), Munich, Germany, June 26-27, 2023*, volume 3456 of *CEUR Workshop Proceedings*, 2023. CEUR-WS.org. URL https://ceur-ws.org/Vol-3456

Book Chapters

66. M. R. Marri, L. Ramachandran, P. Murukannaiah, P. Ravindra, A. Paul, D. Y. Lee, D. Funk, S. Murugappan, and W. HendriX. Dimansionality reduction. In N. F. Samatova, W. Hendrix, J. Jenkins, K. Padmanabhan, and A. Chakraborty, editors, *Practical graph mining with R*, pages 263–310. CRC Press, 2013

Tutorial Notes

67. P. K. Murukannaiah, N. Ajmeri, and M. P. Singh. Ethics in sociotechnical systems, May 2020b. Tutorial notes