Informatics PhD projects at King's College London, AY 24-25 — Human-centred computing

The PhD project proposals listed below will be considered for 2024/25 studentships available in the Department of Informatics to start 1 October 2024 or later during the 2024/25 academic year. Please note that this list is not inclusive and potential applicants can alternatively identify and contact appropriate supervisors outlining their background and research interests or proposing their own project ideas.

The PhD projects are listed in two groups. In the first group are the projects with allocated studentships: each project in this group has one allocated studentship. The remaining studentships will be considered for the projects listed in the second group. The number of those remaining studentships is smaller than the number of the projects in the second group. The allocation of studentships will be based on the merits of individual applications. Applications for PhD studies in the Department of Informatics, for all listed projects as well as for other projects agreed with supervisors, are also welcome from students applying for other funding (within other studentship schemes) and from self-funded students. See also this list of funding opportunities available at King's for post-graduate research in Computer Science.

- Scholarship Allocated
- Scholarship Not Allocated

Scholarship Allocated

(Back to Top)

- Visual live programming in scientific computing and similar domains
- Studying induced demand in software performance
- Inclusive Privacy & Security
- Mitigating Technology-Enabled Harms (Privacy & Security)
- Goal-based explanations for autonomous systems and robots
- Common Sense Planning (for Robotics)
- Adaptation and effective communication in collaborative physically Assistive Tasks
- Explaining robotic planning decision points along execution
- Robotics and Social Justice
- Safe Reinforcement Learning from Human Feedback
- Understanding and Communicating the Ecological Impact of Machine Learning
- Exploring Alternative Permissions Mechanisms for Personal Data Sharing

Visual live programming in scientific computing and similar domains

Supervisor: Stephen Kell

Areas: Systems (SE, programming, autonomous systems, robotics, ...), Visualisation, Data science, Human-centred computing, Computing Applications

(Back to Scholarship Allocated)

Project Description

Currently, working interactively with data means either using a pre-built application offering a fixed interface, which is visual and interactive but offers limited programmability, or using custom workflows built by programming/scripting or command-line wizardry, which are flexible but technically demanding and far less visual and interactive. Computational notebooks like Jupyter are in some senses a third way, being somewhat visual, and have proven approachable by those seeking to learn programming 'on the job'. However, they currently suffer many usability and reproducibility issues, and still present a 'walled garden' environment with poor integration into the surrounding system. This PhD is about ways to combine the interactivity of applications and the flexibility of command lines, possibly by designing a notebook system that works differently than Jupyter et al. We observe that crude operating system (OS) interfaces are the bottleneck to interoperable, visual programming, since they lack a rich data model on which to build visualisation as a system-wide service; this is what leads to smaller-scope walled-garden approaches. Recent work adding run-time type information to native code has shown that such limitations can be overcome without defining an entirely new platform. This project will pursue similar approaches encompassing file data and graphical user interface elements. The objective is to demonstrate a graphical workspace that is highly compatible and interoperable, dealing in files of existing formats, but can support working visually and programmatically via a palette of small, composable, user-tailorable graphical tools. Target audiences include computational scientists, data scientists, digital artists and the like. The project requires systems programming skills and an interest in human--computer interaction topics.

Studying induced demand in software performance

Supervisor: Stephen Kell

Areas: Systems (SE, programming, autonomous systems, robotics, ...), Human-centred computing

(Back to Scholarship Allocated)

Project Description

Tradition has it that computing resources are scarce and therefore that efficiency is a prime concern of programs, especially in infrastructure programs such as compilers. However, in the modern world this is no longer true: computing resources are plentiful and powerful, and software's functionality rarely challenges the limits of the hardware. Despite this, users continue to experience software that is slow, and continue to replace hardware with newer hardware in order to 'keep up' -- especially in the era of continuously updated web-based software. It has long been observed that everyday software is getting slower (e.g. the famous Wirth's law). One theory to explain this is 'induced demand', where greater capacity changes habits of programmers and users in ways that effectively 'soak up' the extra capacity and, often, worsen the apparent infrastructure shortfall. (One classic text on induced demand is Hart & Spivak's "The Elephant in the Bedroom", 1993). This project will study the phenomenon of induced demand in commodity software stacks. It will most likely consist of developing novel profiling tools to study the evolution of the performance of commodity software, and of case studies that pinpoint technical decisions or changes which explain the loss of performance. One possible angle is to study open-source desktop software over the period from the mid-1990s to the present; one tool-building tactic would be to exploit how a single Linux kernel can host user software environments spanning a large interval of time.

Inclusive Privacy & Security

Supervisor: Kovila Coopamootoo

Areas: Cybersecurity, Human-centred computing

(Back to Scholarship Allocated)

Project Description

Privacy, security, digital safety are often designed for majority or WEIRD populations, thereby not appropriately serving all user communities (minority, with specific needs, or life situations), and risk discriminating access and appropriate engagement, while also creating concerns for them. The PhD project will focus on a (type of) privacy, security or digital safety context and seek better privacy/security design with and for at-risk communities, such as addressing:

- Privacy for women and queer communities for e.g. in online platforms
- Authentication for the visually impaired
- Digital safety for socio-economically deprived communities (where privacy/security is often traded-off for more pressing life needs)
- Privacy and digital safety programmes for adult digital starters (often from deprived or minority populations, or experiencing the digital generational divide)
- Privacy and security of refugees and migrants

Proposed PhD projects will involve a mix of qualitative, participatory research and community-centred research. The projects can involve working with stakeholders such as charities, NGOs that support vulnerable and marginalised groups and victim-survivors, as well as digital safety advocacy.

References

- Usenix Security Symposium 2023: "Un-Equal Online Safety?" A Gender Analysis of Security and Privacy Protection Advice and Behaviour Patterns. By Kovila P.L Coopamootoo & Magdalene Ng
- Usenix Security Symposium 2022: "I feel invaded, annoyed, anxious and I may protect myself": Individuals' Feelings about Online Tracking and their Protective Behaviour across Gender and Country. By Kovila P.L Coopamootoo, Maryam Merhnezhad, Ehsan Toreini.
- IEEE Euro S&PW 2023: What we do in the shadows: How does experiencing cybercrime affect response actions and protective practices. By Magdalene Ng, Maria Bada, Kovila P.L Coopamootoo.

- ACM CCS (Computer & Communications Security) 2020: Usage Patterns of Privacy-Enhancing Technologies. By Kovila P.L Coopamootoo.
- Workshop on Privacy in the Electronic Society (WPES) 2014: Sensible privacy: how we can protect domestic violence survivors without facilitating misuse. By Arief, Coopamootoo, Emma, van Moorsel.

Mitigating Technology-Enabled Harms (Privacy & Security)

Supervisor: Kovila Coopamootoo

Areas: Cybersecurity, Human-centred computing

(Back to Scholarship Allocated)

Project Description

There are growing examples of technology originally designed to enhance quality of life being mis-used to facilitate abuse, online harms and in-security. The PhD project will investigate a particular context of technology-enabled harm in intersection with the user characteristic or life event that amplifies the chances of being targeted or in experiencing online harms (e.g. gender, age, life style preference, health condition, breakdown of relationship), such as:

- Online platforms being used for hate and harassment of women, queer communities, or race minorities. This can address how platforms (are required to) mitigate harm and content moderation.
- Smart homes enabling intimate partner abuse
- Digital health and wellbeing technology involving excessive tracking and privacy invasion (e.g. Femtech, sports apps)
- Online platform's poor content moderation and lack of age-appropriate controls that threaten children's online safety
- Content sharing platform used for non-consensual intimate image distribution

Proposed PhD projects will aim to understand the lived experiences of users, make recommendations and prototype for safer/harm-mitigating experiences. Specifically, the project will look into the behavioural dynamics of interacting with such technologies (or human-computer interaction — HCI) and the design aspects, via suitable user-studies such as interviews, surveys, participatory workshops, or ethnography research methods. More information can be found here: https://kovilacoops.github.io/opportunities/

References

- Usenix Security Symposium 2023: "Un-Equal Online Safety?" A Gender Analysis of Security and Privacy Protection Advice and Behaviour Patterns. By Kovila P.L Coopamootoo & Magdalene Ng
- Usenix Security Symposium 2022: "I feel invaded, annoyed, anxious and I may protect myself": Individuals' Feelings about Online Tracking and their Protective Behaviour across Gender and Country. By Kovila P.L Coopamootoo, Maryam Merhnezhad, Ehsan Toreini.

- IEEE Euro S&PW 2023: What we do in the shadows: How does experiencing cybercrime affect response actions and protective practices. By Magdalene Ng, Maria Bada, Kovila P.L Coopamootoo.
- ACM CCS (Computer & Communications Security) 2020: Usage Patterns of Privacy-Enhancing Technologies. By Kovila P.L Coopamootoo.
- Workshop on Privacy in the Electronic Society (WPES) 2014: Sensible privacy: how we can protect domestic violence survivors without facilitating misuse. By Arief, Coopamootoo, Emma, van Moorsel.

Goal-based explanations for autonomous systems and robots

Supervisor: Gerard Canal

Areas: Artificial Intelligence (AI), Human-centred computing

(Back to Scholarship Allocated)

Project Description

Autonomous systems such as robots may become another appliance found in our homes and workplaces. In order to have such systems helping humans to perform their tasks, they must be as autonomous as possible, to prevent becoming a nuisance instead of an aid. Autonomy will require the systems or robots to set up their own agenda (in line with the tasks they are meant to do), defining the next goals to achieve and discarding those that can't be completed. However, this may create misunderstandings with the users around the system, who may expect something different from the robot. Therefore, it is important that these autonomous systems are able to explain why they achieved one task and not another, or why some new (unexpected) task was achieved that was not scheduled. Other sources of misunderstandings may come from action failures and replanning, where the robot finds a new plan to complete an ongoing task. In this case, the new plan may be different to the original one, thus changing the behaviour that the robot was performing. This project will explore how to generate goal-based explanations for robots in assistive/home-based scenarios, extracted from goal-reasoning techniques. It will also look at plan repair to enforce cohesion after a replanning to ideally increase the trust and understanding of the users about the system. Those explanations should also contemplate unforeseen circumstances, therefore explaining things based on "excuses" that the robot may give to the user. Finally, we will investigate how to obtain and provide those explanations at execution time, so explaining on the go. The methods developed shall be integrated into a robotic system, in an assistive/service robot scenario.

References

[1] Canal, G., Borgo, R., Coles, A., Drake, A., Huynh, T. D., Keller, P., Krivic, S., Luff, P., Mahesar, Q-A., Moreau, L., Parsons, S., Patel, M., & Sklar, E. (2020). Building Trust in Human-Machine Partnerships. Computer Law & Security Review, 39.
[2] Hawes, N., Burbridge, C., Jovan, F., Kunze, L., Lacerda, B., Mudrova, L., ... & Hanheide, M. (2017). The strands project: Long-term autonomy in everyday environments. IEEE Robotics & Automation Magazine, 24(3), 146-156.
[3] Aha, D. W. (2018). Goal reasoning: Foundations, emerging applications, and prospects. AI Magazine, 39(2), 3-24.
[4] Bercher, Pascal, et al. "Plan, repair, execute, explain—how planning helps to assemble your home theater." Proceedings of the International Conference on Automated

Planning and Scheduling. Vol. 24. No. 1. 2014.
[5] Chakraborti, Tathagata, Sarath Sreedharan, and Subbarao Kambhampati. "The emerging landscape of explainable AI planning and decision making. IJCAI 2020.
[6] Gobelbecker, M., Keller, T., Eyerich, P., Brenner, M., & Nebel, B. (2010, April). Coming up with good excuses: What to do when no plan can be found. In Proceedings of the International Conference on Automated Planning and Scheduling (Vol. 20, No. 1).

Common Sense Planning (for Robotics)

Supervisor: Gerard Canal / Albert Merono-Penuela

Areas: Artificial Intelligence (AI), Human-centred computing

(Back to Scholarship Allocated)

Project Description

Task Planning (also known as Symbolic Planning or AI Planning) has proved to be a very useful technique to tackle the decision-making problem in robotics. Given a set of task goals, the planner can come up with a set of actions that will reach those goals once executed by the robot. However, plans are often short-lived when robots execute them, given that the real world is complex, and some actions may fail. A traditional approach is to recompute plans once they fail (replanning), however, computing new plans is often costly. This is an issue in robotics and real-time systems, where users wouldn't want a robot that stops for some minutes to compute a plan every now and then. Instead of replanning, a solution could be to repair the plan. While some approaches exist [1, 2], none has yet exploited the semantics of the task and the actions. As the actions are meant to be applied in the real world, the meaning of the action is important and may be used not only to repair plans and post-process them, but also to explain them to users. Moreover, plans involving certain actions and tasks that are not accompanied by a real-world context cannot be guaranteed to be safe or trustworthy for users. While a full specification of task and action semantics is cumbersome due to the size and complexity of open domains, some ongoing efforts are addressing the also general, but more manageable domain of common sense. For example, OpenCyc has been running for decades to "assemble a comprehensive ontology and knowledge base that spans the basic concepts and rules about how the world works" [3]. More recently, the knowledge graph community has advanced ground in integrating various knowledge bases (e.g. ATOMIC, ConceptNet, FrameNet, Roget, Visual Genome, Wikidata, WordNet) of common-sense knowledge, in a hyper-relational graph called Common Sense Knowledge Graph [4] (CSKG). A large number of the symbolic representations (e.g. concepts, relations, rules, etc.) in CSKG are relevant for, and could be used as semantic descriptions of, tasks and actions in robot planning. In this project, we propose to combine the ideas of Symbolic Planning for decision-making in robotics with explicit representations of common-sense knowledge in knowledge graphs for safer planning. The idea is that such a combination can leverage contextual descriptions of domains and use common-sense reasoning to avoid plans containing actions or tasks with the potential of being unsafe or untrustworthy. Furthermore, this may also allow to not only improve plans that might not be executable due to semantic constraints, but also to change them in a way that enhances user trust in the robotic system. The symbolic nature of common-sense knowledge graphs such as the CSKG can provide a layer of explainability ensuring that

plans can be understood and debugged by humans, creating feedback loops between the planner and the knowledge graph. More specifically, this project:

- Assesses common-sense knowledge in explicit symbolic representations, such as those provided by CSKG and other related datasets, as reliable sources of semantic information for robot planning.
- Develops new planning algorithms that leverage common-sense knowledge graphs and common-sense reasoning to propose semantically rich, explainable plans for robots.
- Evaluates the performance, safety and trustworthiness of these implementations by comparing them with existing approaches that do not exploit common sense.

References

[1] Bercher, P., Biundo, S., Geier, T., Hoernle, T., Nothdurft, F., Richter, F., & Schattenberg, B. (2014, May). Plan, repair, execute, explain—how planning helps to assemble your home theater. In Proceedings of the International Conference on Automated Planning and Scheduling (Vol. 24, pp. 386-394).
[2] Fox, M., Gerevini, A., Long, D., & Serina, I. (2006, June). Plan Stability: Replanning versus Plan Repair. In ICAPs (Vol. 6, pp. 212-221).
[3] Lenat, D. B. (1995). CYC: A large-scale investment in knowledge infrastructure. Communications of the ACM, 38(11), 33-38.
[4] Ilievski, F., Szekely, P., & Zhang, B. (2021). Cskg: The commonsense knowledge graph. In The Semantic Web: 18th International Conference, ESWC 2021, Virtual Event, June 6—10, 2021, Proceedings 18 (pp. 680-696). Springer International Publishing.

Adaptation and effective communication in collaborative physically Assistive Tasks

Supervisor: Gerard Canal

Areas: Artificial Intelligence (AI), Human-centred computing

(Back to Scholarship Allocated)

Project Description

Physical robotic Assistance can often be modelled as a collaborative task in which the goal of both the user and the robot is to complete an assistive task together. However, assistive settings have a lot of particularities that differentiate them from traditional Human-Robot Collaboration tasks. For it to be effective, the assistance should be seamless, natural, and without a required effort on the user's side. This means that these robots must be able to communicate with the user in a very natural and intuitive way, but also in an adaptive manner. In this project, we will investigate the development of techniques for the online adaptation of the robot to the human, as well as anticipation of user needs, and seamless communication in the context of assistive tasks such as robotic feeding and dressing.

References

[1] Canal, G., Alenya, G., & Torras, C. (2019). Adapting robot task planning to user preferences: an assistive shoe dressing example. Autonomous Robots, 43(6), 1343-1356.
[2] Canal, G., Alenya, G., & Torras, C. (2016). Personalization framework for adaptive robotic feeding assistance. In Social Robotics: 8th International Conference, ICSR 2016, Kansas City, MO, USA, November 1-3, 2016 Proceedings 8 (pp. 22-31). Springer International Publishing.
[3] Bhattacharjee, T., Lee, G., Song, H., & Srinivasa, S. S. (2019). Towards robotic feeding: Role of haptics in fork-based food manipulation. IEEE Robotics and Automation Letters, 4(2), 1485-1492.
[4] Bhattacharjee, T., Gordon, E. K., Scalise, R., Cabrera, M. E., Caspi, A., Cakmak, M., & Srinivasa, S. S. (2020, March). Is more autonomy always better? exploring preferences of users with mobility impairments in robot-assisted feeding. In Proceedings of the 2020 ACM/IEEE international conference on human-robot interaction (pp. 181-190).
[5] Ondras, J., Anwar, A., Wu, T., Bu, F., Jung, M., Ortiz, J. J., & Bhattacharjee, T. (2022, August). Human-robot commensality: Bite timing prediction for robot-assisted feeding in groups. In 6th Annual Conference on Robot Learning.

Explaining robotic planning decision points along execution

Supervisor: Gerard Canal

Areas: Artificial Intelligence (AI), Human-centred computing

(Back to Scholarship Allocated)

Project Description

Explanation of robotic behaviours has been proved to be very important to improve the understanding of the users of such robots, which improves their trust in the robotic system. However, explanations in robotics are tricky as they need to be given at the correct moment and based on what happened in the execution. In robotic-based planning, an interesting explanation is that of decision points, where the robot could have taken a different action with a different outcome. This project focuses on the explanation of such decision points at execution time, integrating information on current and past events that may help explain the decision to a user. For this, we will look into explainability in the space of plans where, knowing the committed plan and what has happened in the execution, we compare it with the other alternatives that the robot had at a certain decision point. This will evolve towards generating explanations along the execution of plans, as well as determining when some decisions may not be obvious to the user, thus warranting explanations.

References

[1] Canal, G., Torras, C., & Alenya, G. (2023). Generating predicate suggestions based on the space of plans: an example of planning with preferences. User modeling and user-adapted interaction, 33(2), 333-357.
[2] Eifler, R., Cashmore, M., Hoffmann, J., Magazzeni, D., & Steinmetz, M. (2020, April). A new approach to plan-space explanation: Analyzing plan-property dependencies in oversubscription planning. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 34, No. 06, pp. 9818-9826).
[3] Wachowiak, L., Celiktutan, O., Coles, A., & Canal, G. (2023, June). A Survey of Evaluation Methods and Metrics for Explanations in Human—Robot Interaction (HRI). In ICRA2023 Workshop on Explainable Robotics.
[4] Wachowiak, L., Tisnikar, P., Canal, G., Coles, A., Leonetti, M., & Celiktutan, O. (2022, August). Analysing eye gaze patterns during confusion and errors in human—agent collaborations. In 2022 31st IEEE International Conference on Robot and Human Interactive Communication (RO-MAN) (pp. 224-229). IEEE.

Robotics and Social Justice

Supervisor: Martim Brandao

Areas: Artificial Intelligence (AI), Robotics, Human-centred computing

(Back to Scholarship Allocated)

Project Description

The Responsible Robotics and AI Lab is open to applications for a PhD in blue sky research at the intersection of robotics and social justice. The project sits at the intersection of Computer Science and Social Science, and it is expected that the successful candidate will choose an appropriate co-supervisor at a later stage (this choice will be made in conversation with Martim Brandao, the main supervisor). We are particularly looking for students interested in: - Investigating issues of social justice in robotics (e.g. worker conditions, police misuse, accountability, racism, sexism, colonialism) - Developing new methodologies for anti-[racist/ageist/ableist/sexist/capitalist/colonial] robotics - Abolitionist robotics (e.g. for tackling homelessness, mental health, domestic violence, child welfare and other social problems in humane community-grounded ways, without police involvement) - Robotics for sustainability, robotics and environmental justice The PhD proposal should include a plan suggesting how the chosen factors of social justice will be investigated, which (computer simulation-based) prototypes developed, and how they will be evaluated.

Safe Reinforcement Learning from Human Feedback

Supervisor: Yali Du

Areas: Artificial Intelligence (AI), Machine Learning (ML), Human-centred computing

(Back to Scholarship Allocated)

Project Description

Reinforcement learning (RL) has become a new paradigm for solving complex decision making problems. However, it presents numerous safety concerns in real world decision making, such as unsafe exploration, unrealistic reward function, etc. As reinforcement learning agents are frequently evaluated in terms of rewards, it is less noticed that designing AI agents that have the capability to achieve arbitrary objectives can be deficient in that the systems are intrinsically unpredictable and might result in negative and irreversible outcomes to humans. While humans understand the dangers, human involvement in the agent's learning process can be promising to boost AI safety for being more aligned with human values [1]. Dr. Du's early research [2] shows that human preference can be used as an effective replacement for reward signals. One recent attempt [1] also adopted human preference as a replacement for reward signals, to guide the training of agents in safety-critical environments; while agents query humans with a certain probability, how to actively query humans and adapt its knowledge to the task and query is not considered. This project considers to build safe RL agents leveraging human feedback, and aims to address two challenges: 1) how to enable agents to actively query humans with efficiency thus minimising disturbance to humans; 2) how to improve algorithms' robustness in dealing with large state space and even unseen tasks. The target of this project is to realise human value alignment safe RL in a scalable (in terms of task scale) and efficient (in terms of human involvement) way. To address these challenges, this research will leverage the principles of the Abstract Interpretation framework [3], a theory that dictates how to obtain sound, computable, and precise finite approximations of potentially infinite sets of behaviours. Based on the abstraction of states, we aim to enable agents to build a knowledge base for (un)safe behaviours, and thus construct a scheme for when to actively query humans. Besides, due to the nature of sequential decision making, this project will consider temporal abstractions of behaviours and feedback to improve the consistency in safety control. Furthermore, by the effective abstractions, we aim to make the neural-network based agents invariant to task-irrelevant details, and thus generalizable to new downstream tasks.

References

[1] Ilias Kazantzidis, Tim Norman, Yali Du, Christopher Freeman. How to train your agent: Active learning from human preferences and justifications in safety-critical environments. AAMAS 2022.
[2] Runze Liu, Fengshuo Bai, Yali Du, Yaodong Yang. Meta-Reward-Net: Implicitly Differentiable Reward Learning for Preference-based Reinforcement Learning. NeurIPS 2022.
[3] Cousot, P. and Cousot, R. Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In Symposium on Principles of Programming Languages (POPL), 1977.

Understanding and Communicating the Ecological Impact of Machine Learning

Supervisor: Dr Georgia Panagiotidou

Areas: Visualisation, Human-centred computing

(Back to Scholarship Allocated)

Project Description

The training and fine-tuning of Machine Learning (ML) models is costly to the environment in terms of computational power (energy) and ultimately carbon emissions (Bender et al 2021, Schwartz et al 2020, Strubell et al 2019). With the ongoing surge in Machine Learning (ML) in both model numbers as well as their achieved complexity, the environmental impact of ML is adding up to be significant. While the ML and HCI communities has been identifying such critical environmental issues, there is still a lack of comprehensive methodologies to support system designers and developers in taking a sustainability-by-design approach where they understand, measure, and mitigate such issues. For instance, the few existing tools that calculate carbon emission costs (Anthony et al 2020, Lacoste et al 2019) largely focus on technical measurements, largely overlooking if and how they are used by practitioners in the field. This PhD will draw from research on sustainable HCI, behaviour-change and critical data visualisation to map out the state of the field and examine how to best communicate ML's potential environmental effects. Specifically the student will investigate current sustainability-oriented practices within ML teams and strive towards developing a design toolkit for supporting more sustainable-oriented decision-making. The project may include activities such as conducting interviews and observations, running co-design workshops as well as developing visualisations of carbon emissions of machine learning models to act as a tools meant for eco-feedback.

References

Bender E. M., Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21). Association for Computing Machinery, New York, NY, USA, 610—623. https://doi.org/10.1145/3442188.3445922

Costanza-Chock, Sasha. Design justice: Community-led practices to build the worlds we need. The MIT Press, 2020.

D'ignazio, Catherine, and Lauren F. Klein. Data feminism. MIT press, 2023.

Panagiotidou G., Enrico Costanza, Michael J. Fell, Farhan Samanani, and Hannah Knox. 2023. Supporting Solar Energy Coordination among Communities. Proc. ACM Interact.

Mob. Wearable Ubiquitous Technol. 7, 2, Article 71 (June 2023), 23 pages. https://doi.org/10.1145/3596243
Schwartz R., Jesse Dodge, Noah A. Smith, and Oren Etzioni. 2020. Green AI. Commun. ACM 63, 12 (December 2020), 54— 63. https://doi.org/10.1145/3381831
Strubell E., Ananya Ganesh, and Andrew McCallum. 2019. Energy and Policy Considerations for Deep Learning in NLP.
In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, pages 3645—3650, Florence, Italy. Association for Computational Linguistics.

Exploring Alternative Permissions Mechanisms for Personal Data Sharing

Supervisor: William Seymour (co-supervisor TBC)

Areas: Cybersecurity, Human-centred computing

(Back to Scholarship Allocated)

Project Description

Installing a new app or unpacking a new smart home device almost always involves granting a range of permissions about how those products can use your personal data. Browsing the web similarly requires making a vast number of decisions about cookies and other tracking technologies. But giving consent in this way often doesn't really feel like consenting at all, and prior work in this area has highlighted the implausibility that the mechanisms we use fulfil the requirements of what is understood as informed consent. This is particularly apparent in contexts where multiple people are using the same device. "Bystander" partners, children, family, and housemates are often left out of the installation process, even though information about them is also being collected. Instead, one member of the household unilaterally makes decisions for everybody else. This echoes the model adopted by data protection regulation, which (broadly) controls the use of data on those within the home by organisations outside of it, avoiding the topic of privacy between cohabitants. This PhD project will build on my and other existing work in this area to explore alternatives to the status quo described above, including the potential for novel automated and group privacy decision making mechanisms. It will do so across a range of different interaction modalities (such as graphical, voice, and touch interfaces) and social contexts. You'll be able to shape the direction of the project to fit your own interests. Key skills and research methodologies will include some of the following, tailored to your preferences: -Qualitative interviews and/or focus groups with users and designers -Quantitative surveys with users -Lightweight prototyping and speculative design (potentially including programming prototypes) -More creative methods such as home deployments and creating interactive experiences While you don't need to have experience with any of these techniques before starting, you do need to have a demonstrable interest/experience in online privacy and human computer interaction.

References

William Seymour, Mark Cote, and Jose Such. 2023. Legal Obligation and Ethical Best Practice: Towards Meaningful Verbal Consent for Voice Assistants. In Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 166, 1—16.

https://doi.org/10.1145/3544548.3580967
William Seymour, Reuben Binns, Petr Slovak, Max Van Kleek, and Nigel Shadbolt. 2020. Strangers in the Room: Unpacking Perceptions of 'Smartness' and Related Ethical Concerns in the Home. In Proceedings of the 2020 ACM Designing Interactive Systems Conference (DIS '20). Association for Computing Machinery, New York, NY, USA, 841—854. https://doi.org/10.1145/3357236.3395501
Bettina Nissen, Victoria Neumann, Mateusz Mikusz, Rory Gianni, Sarah Clinch, Chris Speed, and Nigel Davies. 2019. Should I Agree? Delegating Consent Decisions Beyond the Individual. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19). Association for Computing Machinery, New York, NY, USA, Paper 515, 1—13. https://doi.org/10.1145/3290605.3300745

Scholarship Not Allocated

(Back to Top)

- Understanding the Complexity of Negotiations
- Designing with and for at-risk populations
- Female mobile health apps: understanding user needs and concerns
- Privacy, safety, and trust in emerging AI technologies

Understanding the Complexity of Negotiations

Supervisor: Dr Alfie Abdul-Rahman & Dr Rita Borgo

Areas: Visualisation, Human-centred computing, Natural Language Processing (NLP), Social computing, Machine Learning (ML)

(Back to Scholarship Not Allocated)

Project Description

A negotiated text is the product of a formal decision-making process where a text has been negotiated and drafted over a period of time. Many of the foundational texts of the modern world have not been written by individuals, by negotiated by groups of people in formal settings. For example, treaties between states such as the Universal Declaration of Human Rights or the Treaty of Versailles; or constitutions, such as the one negotiated by the American states in the Constitutional Convention of 1787. During such negotiations, it is important for us to keep track of the delegations and their involvements to grasp their influence on the negotiation process either using techniques such as close reading, distance reading, or machine learning. Even relatively short historical documents written collectively in this way have been the product of thousands of specific proposals and decisions. This project will apply a visual analytics approach towards the understanding of the complexity of a negotiation and the influence of the delegations during a negotiation process. Possible research questions: a. Developing new static and interactive visualization to assist with data discovery and insight generation in large datasets of events within interacting timelines. b. Developing new approaches to show the evolution of complicated, technical documents over the period of months or years. c. Developing new approaches for indexing the datasets related to the negotiation of documents, and more intuitive displays of the results. d. Developing natural-language-based approaches to relate information captured in 'informal' archives (such as private diaries, letters, social media feeds etc.) to the formal records of a negotiation. This project will work closely with the Quill Project, based at Oxford University:

References

1. Exploring Interpersonal Relationships in Historical Voting Records. G. D. Cantareira, Y. Xing, N. Cole, R. Borgo, and A. Abdul-Rahman. In Computer Graphics Forum, 42(3):211-221, 2023.
2. Quill: A Framework for Constructing Negotiated Texts — with a Case Study on the US Constitutional Convention of 1787. N. Cole, A. Abdul-Rahman, and G. Mallon. In ACM/IEEE Joint Conference on Digital Libraries (JCDL), 2017.

Designing with and for at-risk populations

Supervisor: Ruba Abu-Salma

Areas: Cybersecurity, Human-centred computing, Social computing

(Back to Scholarship Not Allocated)

Project Description

Recently, researchers have started to realize that designing digital technologies for one population in mind risks ignoring the security, privacy, and safety needs of, as well as creating concerns for, other populations. While research on understanding the needs and concerns of at-risk populations like older adults or migrant domestic workers, is evolving with the US dominating the field, the research is still in its infancy without a clear understanding of the scale and impact of technologies on such populations. The objective of this project is to empirically study the security, privacy, and safety needs — based on lived experiences — of at-risk populations like the poor, the young, and the disabled using qualitative (eg, interviews, focus groups, participatory design workshops) and quantitative (eg, surveys) methods. The empirical evidence gleaned from these studies will inform the design of current and future technologies.

References

- https://arxiv.org/pdf/2112.07047.pdf.
- https://arxiv.org/pdf/2206.15037.pdf.
- https://www.usenix.org/conference/soups2022/presentation/bernd.
- https://www.usenix.org/conference/foci20/presentation/bernd.
- https://www.usenix.org/conference/usenixsecurity22/presentation/slupska-vulnerability.
- https://petsymposium.org/popets/2024/popets-2024-0009.pdf.

Female mobile health apps: understanding user needs and concerns

Supervisor: Ruba Abu-Salma

Areas: Cybersecurity, Social computing, Human-centred computing

(Back to Scholarship Not Allocated)

Project Description

Mobile apps which support women's health have developed rapidly. However, the ubiquity of these apps has advanced the practice of intimate surveillance and sensitive data collection. While the overturning of Roe v. Wade has prompted reflection on the privacy and safety implications of female mobile health apps, the needs and concerns of the users of these apps are yet to be explored. The project aims to generate empirical evidence of these needs and concerns, as well as improve the design of female health apps, with a focus on user privacy and safety.

References

- https://petsymposium.org/popets/2020/popets-2020-0083.pdf.
- https://arxiv.org/pdf/2306.05956.pdf.
- https://dl.acm.org/doi/pdf/10.1145/3411764.3445132.

Privacy, safety, and trust in emerging AI technologies

Supervisor: Ruba Abu-Salma

Areas: Cybersecurity, Social computing, Human-centred computing

(Back to Scholarship Not Allocated)

Project Description

Research examining the impact of AI on society is concerned with the need to develop AI technologies that are ethical, human-centered, responsible, safe, and trustworthy. However, the body of human-subjects research investigating how people may define those goals is not large, and little of it compares how views vary across countries and populations; eg, with respect to demographics, social norms, technological experiences, etc. The proposed research aims to expand our empirical understanding of people's experiences with and views on AI technologies, including GenAI tools, across different countries and populations. In particular, we are interested in how privacy and safety concerns affect people's trust in these technologies.

References

- https://arxiv.org/pdf/2302.05284.pdf.
- https://arxiv.org/pdf/2305.06415.pdf.
- https://arxiv.org/pdf/2302.08157.pdf.
- https://arxiv.org/pdf/2310.06778.pdf.
- https://dl.acm.org/doi/pdf/10.1145/3613249.
- https://www.usenix.org/system/files/soups2023-kelley.pdf.