

# Artificial Intelligence for Collective Intelligence: A National-Scale Research Strategy

SETH BULLOCK<sup>1\*</sup>, NIRAV AJMERI<sup>1</sup>, MIKE BATTY<sup>2</sup>, MICHAELA BLACK<sup>3</sup>, JOHN CARTLIDGE<sup>4</sup>, ROBERT CHALLEN<sup>4</sup>, CANGXIONG CHEN<sup>5</sup>, JING CHEN<sup>6</sup>, JOAN CONDELL<sup>3</sup>, LEON DANON<sup>4</sup>, ADAM DENNETT<sup>2</sup>, ALISON HEPPENSTALL<sup>7</sup>, PAUL MARSHALL<sup>1</sup>, PHIL MORGAN<sup>8</sup>, AISLING O’KANE<sup>1</sup>, LAURA G. E. SMITH<sup>9</sup>, THERESA SMITH<sup>10</sup> and HYWEL T. P. WILLIAMS<sup>11</sup>

<sup>1</sup>*School of Computer Science, University of Bristol*; <sup>2</sup>*Centre for Advanced Spatial Analysis, University College London*; <sup>3</sup>*School of Computing, Engineering & Intelligent Systems, Ulster University*; <sup>4</sup>*School of Engineering Mathematics and Technology, University of Bristol*; <sup>5</sup>*Institute for Mathematical Innovation, University of Bath*; <sup>6</sup>*School of Mathematics, Cardiff University*; <sup>7</sup>*School of Social and Political Sciences, University of Glasgow*; <sup>8</sup>*School of Psychology, Cardiff University*; <sup>9</sup>*Department of Psychology, University of Bath*; <sup>10</sup>*Department of Mathematical Sciences, University of Bath*; <sup>11</sup>*Department of Computer Science, University of Exeter*  
\*E-mail: seth.bullock@bristol.ac.uk

## Abstract

Advances in artificial intelligence (AI) have great potential to help address societal challenges that are both *collective* in nature and present at *national* or trans-national scale. Pressing challenges in healthcare, finance, infrastructure and sustainability, for instance, might all be productively addressed by leveraging and amplifying AI for national-scale *collective intelligence*. The development and deployment of this kind of AI faces distinctive challenges, both technical and socio-technical. Here, a research strategy for mobilising inter-disciplinary research to address these challenges is detailed and some of the key issues that must be faced are outlined.

## 1 Introduction

Artificial intelligence (AI) and machine learning often address challenges that are relatively monolithic: determine the safest action for an autonomous car; translate a document from English to French; analyse a medical image to detect a cancer; answer a question about a difficult topic. These kinds of challenge are important and worthwhile targets for AI research. However, an alternative set of challenges exist that are *collective* in nature:

- help to minimise a pandemic’s impact by coordinating mitigating interventions;
- help to manage an extreme weather event using real-time physical and social data streams;
- help to avoid a stock market crash by managing interactions between trading agents;
- help to guide city developers towards more sustainable coordinated city planning decisions;
- help people with diabetes to collaboratively manage their condition while preserving privacy.

The capability of naturally occurring collective systems to solve problems of coordination, collaboration, and communication has been a long-standing inspiration for engineering (Bonabeau et al. 1999). However, developing AI systems for these types of problem presents unique challenges: extracting reliable and informative patterns from multiple overlapping and interacting real-time data streams; identifying and controlling for evolving community structure within the collective; determining local interventions that allow smart agents to influence collective systems in a positive way; developing privacy-preserving machine learning; advancing ethical best practice

and governance; embedding novel machine learning and AI in portals, devices and tools that can be used transparently and productively by different types of user. Tackling them demands moving beyond typical AI/machine learning approaches to achieve an understanding of relevant group dynamics, collective decision-making and the emergent properties of multi-agent systems, topics more commonly studied within the growing research area of collective intelligence. Consequently, addressing these challenges requires a productive *combination* of collective intelligence research and artificial intelligence research (Berditchevskaia & Baeck 2020, Berditchevskaia et al. 2022).

In this paper we introduce and detail a research strategy for approaching this challenge that is being taken by a new national artificial intelligence research hub for the United Kingdom: *AI for Collective Intelligence (AI4CI)*.<sup>1</sup> The AI4CI Hub is a multi-institution collaboration involving seven partner universities from across the UK’s four constituent nations and over forty initial stakeholder partners from academia, government, charities and industry. It pursues applied research at the interface between the fields of AI and collective intelligence, and works to build capacity, capability and community in this area of research across the UK and beyond. This paper presents the AI4CI research strategy, details how it can be pursued across multiple different research themes, and summarises some of the key unifying research challenges that it must address.

## 2 Research Context

Between 2022 and 2024, the UK government initiated several significant investments in national-scale AI research amounting to approximately £1Bn of support. Foremost amongst these investments were: the establishment of the UK’s first national supercomputing facility for AI research (Isambard-AI; £225m),<sup>2</sup> plus an additional £500m of AI compute hardware investment across UK universities,<sup>3</sup> the inception of twelve new Centres for Doctoral Training in AI (AI CDTs; £117m),<sup>4</sup> funding for a raft of AI research projects including AI for net zero (£13m)<sup>5</sup> and AI for healthcare (£13m),<sup>6</sup> the creation of UK Responsible AI, a national network to conduct and fund research into responsible AI (UKRAI, £31m),<sup>7</sup> and the launch of nine new national Research Hubs for AI, three focusing on the mathematical foundations of AI and six focusing on applied AI research (£100m).<sup>8</sup>

These significant investments were driven by growing recognition that modern AI has the potential to achieve a positive and revolutionary impact on society. Here, we focus on a research strategy proposed in response to the findings of a recent Nesta report<sup>9</sup> which recommended that policymakers “put collective intelligence at the core of all AI policy in the United Kingdom” (Berditchevskaia & Baeck 2020, p.57), arguing that “the first major funder to put £10 million into this field will make a lasting impact on the future trajectory for AI and create new opportunities for stimulating economic growth as well as more responsible and democratic AI development” (ibid, p.58).

<sup>1</sup><https://ai4ci.ac.uk>

<sup>2</sup><https://www.gov.uk/government/news/bristol-set-to-host-uks-most-powerful-supercomputer-to-turbocharge-ai-innovation>

<sup>3</sup><https://www.ukri.org/opportunity/host-sites-for-the-next-wave-of-uk-government-ai-infrastructure>; N.B. this investment in hardware was subsequently withdrawn by the UK’s incoming Labour government: <https://www.bbc.co.uk/news/articles/cyx5x44vnyeo>

<sup>4</sup><https://www.ukri.org/what-we-do/how-we-work-in-ai/ukri-artificial-intelligence-centres-for-doctoral-training>

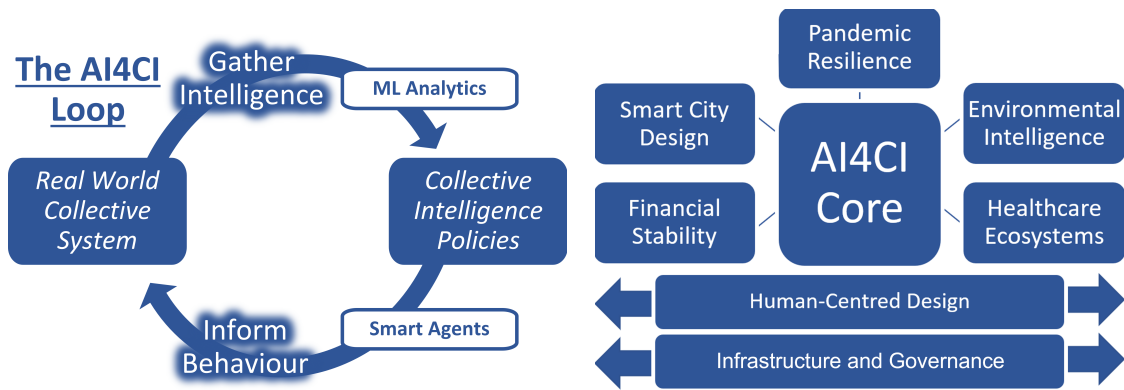
<sup>5</sup><https://www.ukri.org/opportunity/artificial-intelligence-research-to-enable-uks-net-zero-target>

<sup>6</sup><https://www.ukri.org/news/13-million-for-22-ai-for-health-research-projects>

<sup>7</sup><https://www.ukri.org/news/54m-to-develop-secure-ai-that-can-help-solve-major-challenges>

<sup>8</sup><https://www.ukri.org/news/100m-boost-in-ai-research-will-propel-transformative-innovations>

<sup>9</sup><https://www.nesta.org.uk/report/future-minds-and-machines>



**Figure 1** *Left* — The AI4CI Loop: Machine learning and AI enable distributed real-time data streams to inform effective collective action via smart agents. *Right* — The AI4CI Hub: Five applied research themes and two cross-cutting research themes are supported by the hub’s central core.

### 3 Vision and Structure

Our ability to address the most pressing current societal challenges (e.g., healthcare, sustainability, climate change, financial stability) increasingly depends upon the extent to which we can reliably and successfully engineer important kinds of collective intelligence, which we define as:

“Connected communities of people, devices, data and software collaboratively sensing and interacting in real time to achieve positive outcomes at multiple scales”.

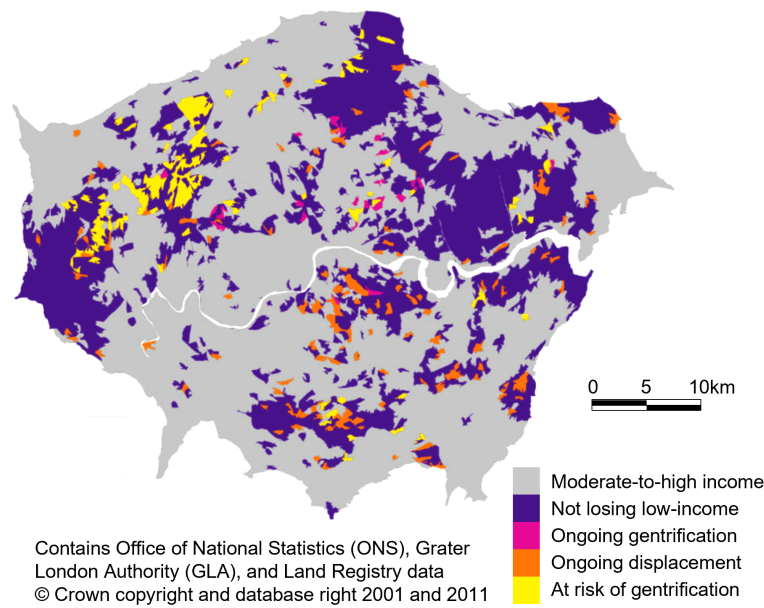
Whether we are aiming to minimise the impact of a global pandemic through effectively managing successive waves of vaccination (Brooks-Pollock et al. 2021), to prevent financial “flash crashes” through effective regulation of autonomous trading agents (Cartlidge et al. 2012, Johnson et al. 2013), to make our cities sustainable and liveable through using real-time analytics to inform short-, medium- and long-term planning (Spooner et al. 2021, Batty 2024), to combat social polarisation and climate disinformation on social media (Treen et al. 2020), or to achieve the UK NHS 2019 Long-Term Plan (NHS 2019) by enabling effective healthcare ecosystems that integrate clinical care, technology, education, and social support for patients with chronic health conditions (Duckworth et al. 2024), invariably what is required is an ability to engineer *smart collectives*.<sup>10</sup>

This will necessarily involve addressing both halves of what we characterise as the “AI4CI Loop” (Fig. 1 *left*) — (1) *Gathering Intelligence*: collecting and making sense of distributed information; (2) *Informing Behaviour*: acting on that intelligence to effectively support decision making at multiple levels. New artificial intelligence methods are unlocking progress on both halves. The first is being revolutionised by a combination of mobile devices, instrumented environments, data science, machine learning analytics and real-time visualisation, while the second is being transformed by decentralised, distributed smart agent technologies that interact directly with users. However, successfully linking both halves of the AI4CI Loop requires new human-centred design principles, new governance practices and new infrastructure appropriate for systems that deploy AI for collective intelligence at national scale.

### 4 Research Themes

The AI for Collective Intelligence Hub (Fig. 1 *right*) addresses and connects both halves of the AI4CI Loop across a set of five important application domains (healthcare, finance, the environment, pandemics, and cities) and two cross-cutting themes (human-centred design and

<sup>10</sup>This characterisation of collective intelligence is strongly aligned with approaches developed within socio-technical systems research (Baxter & Sommerville 2011).



**Figure 2** An indicative snapshot of smart city datasets informing AI for collective intelligence research. Gentrification and displacement typologies for Greater London in 2011 at neighbourhood level with cartogram distortion based on London’s residential population in 2011. Adapted from Zhang et al. (2020).

infrastructure and governance). In each domain, the challenge is to leverage and make sense of real time, dynamic data streams generated across hybrid systems of interacting people, machines and software distributed over space and across networks, in order to achieve systemic insights and drive effective interventions via the automated behaviour of smart AI agents. Pursuing research across multiple domains in concert enables each to benefit from the others’ insights and maximises the chance of uncovering principles and that have domain-general application (Smaldino & O’Connor 2022).

#### 4.1 Smart City Design

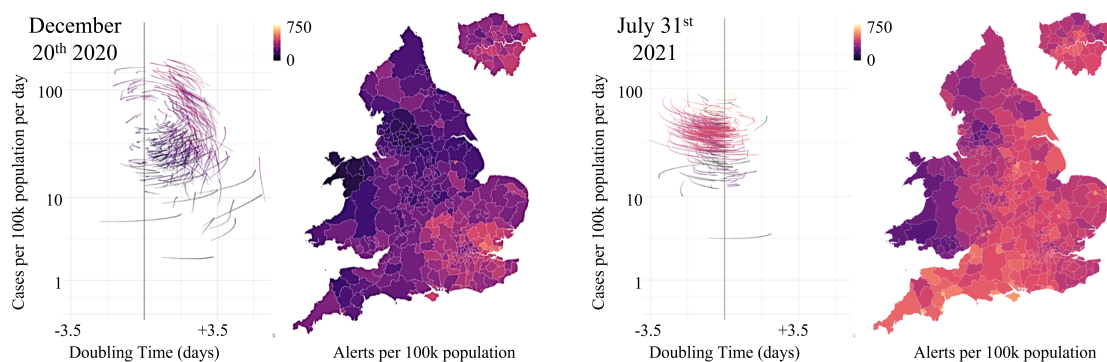
Plan-making systems for UK cities are not currently fit for purpose.<sup>11</sup> Local plans, the major instrument of the statutory planning system, must be modernised to exploit collective data and machine intelligence (Batty 2024). Meeting the challenges associated with smart planning for smart cities<sup>12</sup> in a way that delivers practical tools and applications requires integrating and exploiting multiple streams of city data provided by local and national government, urban analytics and infrastructure firms, national agencies, survey data, and human mobility patterns derived from digital traces or social media.<sup>13</sup>

These data can drive new AI for two purposes: (i) automating real-time intelligence for the smart city (Malleon et al. 2022, An et al. 2023); (ii) informing longer-term smart city planning to

<sup>11</sup>“Planning for the future”, Department for Levelling Up, Housing and Communities, UK Government, 2023, <https://www.gov.uk/government/consultations/planning-for-the-future/planning-for-the-future>

<sup>12</sup>The term *smart city* is used here in two mutually reinforcing senses, first in the sense that novel smart technologies are physically incorporated into these cities, and second in the sense that these technologies underpin new kinds of “smart” behavioural interactions within and across these cities at a range of different time scales (Batty et al. 2020).

<sup>13</sup>e.g., Data for Good: <https://dataforgood.facebook.com>; Smart Data Research: <https://www.sdruk.ukri.org> (formerly Digital Footprints); note the implicit challenges here related to (i) establishing and maintaining the public’s trust in, and engagement with, these potentially intrusive data collection efforts, and (ii) countering the inevitable systematic biases that arise from unrepresentative sampling of the collective system as a whole.



**Figure 3** A snapshot of pandemic datasets informing AI for collective intelligence research. Regionally disaggregated datasets relate the level and growth rate of COVID-19 cases (phase plots) with the rate of digital contact tracing alerts delivered to citizens by the NHS mobile phone app (maps) at two points in time during the COVID-19 pandemic. *Left* — December 20<sup>th</sup> 2020: the alpha variant is spreading in the south-east despite a ‘circuit-breaker’ lockdown. *Right* — July 31<sup>st</sup> 2021: Digital contact tracing alerts are triggered by high COVID-19 case burden

meet the challenges of climate, ageing, housing affordability, and health (Batty 2024). Achieving smarter cities that optimise behaviours in the short and long term requires AI that extends and improves on existing models of urban structure, dealing with highly fluid situations dominated by rapid change (Batty 2024). This is a major challenge not only for the way that we design cities but also for how AI must deal with many/most human problem-solving contexts. Supervised and unsupervised learning methods can be used to reveal new patterns in large messy mobility datasets such as mobile phone traces, cross-validated with rich survey data to produce spatially, temporally and attribute rich insights into the seismic shift in post-COVID mobility patterns (Batty et al. 2020). Predictive tools for the design of new patterns of transport and land development at different scales can be founded on models that take multiple land suitability and mobility indices as inputs (see, e.g., Fig. 2; Zhang et al. 2020). Deriving meaningful interpretations of these models and enabling decision-makers to explore how optimal plans play out over time and space in the context of synthetic AI agent models (Batty et al. 2012) delivers the explanatory accounts that are essential for public accountability in the use of these methods for city decision making.

#### 4.2 Pandemic Resilience

COVID-19 exposed weaknesses in the UK’s pandemic resilience. A combination of collective intelligence and AI can help us do better next time. Data crucial for managing novel pandemics are inherently fragmented, arising from communities of medics, public health professionals and analysts to describe the spread of the disease, characterise its phenotype, and, with appropriate modelling, inform appropriate policies nationally and locally (Brooks-Pollock et al. 2021). National spatio-temporal datasets describing SARS-CoV-2 hospital testing and community testing, and the extent and effects of the mitigations put in place against it, can be exploited in order to build new AI/machine learning tools for future pandemics—and potentially for mitigating seasonal outbreaks of endemic disease.

Two strands of research can be identified. First, a suite of machine learning models fuelled by national SARS-CoV-2 pandemic data can be used to explore and demonstrate how the integration of multiple population-level indicators could have improved decision making during the pandemic. Due to the urgency of the need for response during the pandemic, developing and validating this kind of analytical infrastructure was not possible. With it in place, however, challenges associated with imperfect data can be addressed. For example, data from unreliable laboratories distorted

local and national COVID-19 reproduction number ( $R_t$ ) estimates in 2022<sup>14</sup> leading to under-informed policy decisions. Automatic detection and correction for such errors will alert policy makers and decision makers more quickly.

Second, detailed local data can be used to understand localised interventions and spontaneous behavioural responses. Here, relevant AI challenges include: establishing the relative value of diverse data streams at different stages of the pandemic and their consistency across spatial scales; the use of anomaly and change point detection to identify meaningful discontinuities; robust data imputation, pattern completion, bias detection and correction; evaluating the impact of both vaccination and behavioural change resulting from, for example, non-pharmaceutical interventions (and their interactions); and coping with delay and bias in data capture and clinical outcome results during the exponential growth phase of a new variant. One particular focus is the impact of contact tracing apps on the behaviour of individuals and the public-health messaging around their use (see Fig. 3).

Two major challenges cut across both strands: quality assurance within a privacy protecting framework (Challen et al. 2019) and managing the regional heterogeneity of pandemic impact and associated behavioural change (Challen et al. 2021).

If these can be overcome, there is the potential to deliver a suite of tools (sensitive to local population properties: income, mobility, demographics) that inform policy on where and when to increase or decrease testing capacity, implement contact tracing and/or isolation measures, distribute limited hospital capacity, etc. By working in collaboration with key stakeholders in government, a set of interactive portals can be developed that effectively inform policy decisions (by reporting well understood metrics: the reproduction number,  $R_t$ , hospitalisations, excess deaths) and/or individual behaviour (e.g., by presenting bespoke risk scenarios to increase adherence to government-imposed restrictions).

### 4.3 Environmental Intelligence

In order to meet the challenge of mitigating climate change, there is a need to improve access to, and comprehension of, different kinds of complex, time-varying environmental data, for example: (1) outputs from climate, weather and ocean model ensembles and empirical observations, (2) high-volume geospatial, ecological, satellite and remote sensing data; (3) socioeconomic data on resource flows, supply chains, energy consumption and carbon emissions, and (4) online media including news media and cross-platform social media content. Many decision-makers (including citizens and policy-makers) would benefit from better environmental information. A huge volume of this data is now available, but it often requires a high level of expertise to obtain it and interpret the associated uncertainties. Large language models are increasingly suggested as a potential solution to part of this problem (Vaghefi et al. 2023, Koldunov & Jung 2024). Meanwhile, public debate is weakened by the profusion of poor quality or deliberately false information, especially concerning the contested issue of climate change (Treen et al. 2020, Acar 2023). A combination of AI and collective intelligence approaches can deliver tools that overcome these challenges by democratising access to good quality information about environmental change.

Novel “climate avatar” agents can act as simple interfaces between complex environmental data and the people who need it in order to improve their decision-making. They ingest weather and climate data from existing large datasets,<sup>15</sup> peer-reviewed climate science literature and other trusted sources (e.g., IPCC reports<sup>16</sup>), and expose this data via natural-language interfaces that allow users to access information and gain understanding in a conversational style. By summarising scientific literature and generating on-the-fly visualisations from raw data, climate

<sup>14</sup><https://www.gov.uk/government/news/ukhsa-publishes-investigation-findings-following-errors-at-the-private-immensa-lab>

<sup>15</sup>e.g., <https://www.metoffice.gov.uk/research/approach/collaboration/ukcp>, the UK Met Office’s Climate Projections dataset.

<sup>16</sup><https://www.ipcc.ch/reports/>

avatars enable lay users to make sense of complex climate data and uncertainties, tailored to their specific context (e.g., where they live or the sector in which they work). Similar smart agents engage with other environmental data sources, such as geospatial data, satellite imagery and ecological data. Once created, validated and trusted, these avatars can be deployed to interact with human users in different contexts, for example: allowing expert and non-expert academics to interrogate complex federated models of natural and human capital; enabling chatbots to explain extreme weather events and provide warnings/guidance; providing timely responses to policy formulation queries; or defusing toxic discourse on social media (Treen et al. 2020). Data ethics, governance and usability challenges must be addressed in order to ensure that agents are explainable, trustworthy, and able to effectively influence public understanding in order to achieve positive social outcomes.

#### 4.4 Financial Stability

Modern financial technology (FinTech) presents major challenges for both regulation (cf. the UK Government’s 2021 Kalifa Review<sup>17</sup>) and consumer protection/trust (cf. the UK Financial Conduct Authority’s 2022 Consumer Duty<sup>18</sup>). These challenges can be addressed through collaboration with relevant SMEs, non-profits, consultancies, national research organisations, data providers, platform providers and national FinTech hubs, to co-create personalised AI for early warning indicators, informing financial decision making, and reducing vulnerability to manipulation.

The development of personalised, adaptive AI driven by collective intelligence derived from financial systems to enable individuals, businesses and government to make better-informed decisions can be considered in two contexts: (i) financial markets (Shi & Cartlidge 2022, Buckle et al. 2023, Liu et al. 2023, Zhang, Wen & Chen 2023, Shi & Cartlidge 2024, You et al. 2024) and (ii) personal finances (Bazarbash 2019, van Thiel & Elliott 2024). Progress on either requires that systems take advantage of a range of rich and often non-traditional financial datasets: e.g., high frequency financial trading data, equity investment data for retail traders, personal finance, lending records and psychometric credit ratings, demographic household data, NLP-enhanced social media and news, and FCA-approved synthetic datasets used for testing all aspects of financial technology for the UK’s Financial Conduct Authority regulatory sandbox.

Models that leverage social media, news media and traditional data (prices, volumes, etc.) can provide regulators with early warning signals of bubbles/crashes and protect non-professional investors from pathological investment behaviour (e.g., the viral herding that led to the recent GameStop short squeeze; Klein 2022, Dambanemuya et al. 2023) with model validation taking place within the Financial Conduct Authority sandbox.<sup>19</sup> AI assistants that enable collective ethical investment, and protect vulnerable households from exploitative personal finance providers can be co-created with relevant charities, e.g., working with data-driven psychometric credit rating tools that personalise the immediate and longer-term implications of personal and SME borrowing decisions. Overall, work in this area can make use of rich forms of non-traditional financial data to develop new AI tools that can de-risk the rapid “democratisation of finance” that is being provoked by the ongoing FinTech revolution (Arner et al. 2015).

#### 4.5 Healthcare Ecosystems

The ability of the UK’s NHS to deliver its Long-Term Plan depends critically on its capacity to automate bespoke monitoring and support for a range of long-term health conditions at population scale (Topol 2019). This cannot be achieved without a step change in the use of longitudinal data analysis and smart software assistants. Doing so will require working in close

<sup>17</sup><https://www.gov.uk/government/publications/the-kalifa-review-of-uk-fintech>

<sup>18</sup><https://www.fca.org.uk/firms/consumer-duty>

<sup>19</sup><https://www.fca.org.uk/firms/innovation/digital-sandbox>

collaboration with clinical partners and technology firms on healthcare analytics and the design of user experience for healthcare AI.

Anonymised patient records that track, e.g., mental health consultations or diabetes progression are complex, partial and noisy reflections of longitudinal patient trajectories with potential to improve clinical decision making and empower patients (Rajkomar et al. 2019). The UK’s NHS trusts and healthcare technology firms have extensive expertise in leveraging data to manage and treat these conditions, including cohorts of diabetes patients engaged in the co-design of AI systems for collective intelligence that are trustworthy and effective (Duckworth et al. 2024).

Here, two interacting strands of research can be identified: (i) machine learning analytics for collective healthcare data, and (ii) co-design of smart healthcare agents for patient collectives. Strand (i) develops methods for unsupervised extraction and quantification of patterns from patient data pooled across heterogeneous sources from clinical systems to networked personal devices in order to discover clusters in symptom trajectories, detect adverse events, and recommend treatment and self-care strategies. This work leverages cutting-edge privacy preserving and federated machine learning methodologies to enable machine learning on data across all the sources interactively and in real-time while guaranteeing that the identity of the individual patients and the data they provide will not be leaked through for example training-data leakage attacks (Chen & Campbell 2022). Strand (ii) works with hard-to-reach patients (those suffering from secondary health conditions, mental health conditions, or living circumstances that prevent them from accessing health care unaided and limit their use of technology), plus their carers and clinicians, to address issues of trust, usability, and efficacy in ethical AI for informing healthcare decision making across patient populations (Stawarz et al. 2023). The over-arching challenge for both strands is to leverage population-wide data collection for informing robust individualised decision-making without compromising anonymity and under realistic data and user assumptions. A key challenge is using AI to mediate between patients and care systems rather than burdening already overloaded clinicians with another software tool (Emanuel & Wachter 2019).

## 5 Cross-cutting Themes

The application domains described above are by no means the only areas in which AI for collective intelligence has strong potential. Additional problems for which productive work could have significant transformative effects include preventing violent extremism (Smith et al. 2020, Bullock & Sayama 2023), addressing the climate crisis (Góis et al. 2019), collaborating with autonomous systems (Pitonakova et al. 2018, Hart et al. 2022), and reducing energy consumption (Bourazeri & Pitt 2018). However, in addition to confronting issues specific to each of these individual use cases, achieving AI for collective intelligence also faces challenges that cut across these application areas.

### 5.1 Human-Centred Design

One issue vital to developing AI for collective intelligence within any use domain is achieving successful interaction with human users. Methods from social and cognitive psychology and human factors must be integrated with the various kinds of research activity outlined above in order to derive human-centred design principles for effective, trustworthy AI agents that inform behavioural change at scale within socio-technical human-AI collectives.

Three parallel strands can be identified: (i) bringing human-centred design considerations to the work within various domain-specific AI for collective intelligence research themes; (ii) developing usable smart agents that assist users in accessing, understanding, and acting on guidance derived from collective intelligence data; and (iii) pursuing fundamental questions related to understanding and managing “tipping points” in collective intelligence systems. For (i), participatory design methods (Bratteteig et al. 2012) involving academics and stakeholders can be employed to prototype human-machine interfaces (HMIs) for the AI systems being developed.



For (ii), data from human experiments can inform a series of design iterations, focussing on accessibility, usability, explainability, adaptability and trust (Choung et al. 2023), drawing upon long-standing approaches to defining and measuring trust in automation (e.g., Lee & See 2004), with all being key factors for the acceptance, adoption and continued use of new technologies. Comparative analyses of these data reveal transfer effects between different theme settings, guiding development of demonstrators within each domain. For (iii), testable predictions of how to identify, characterise and influence tipping point thresholds for behaviour change can be derived from data on explainability, confidence, persistent adoption, praise and blame, e.g., based on the perceived capability of the system (Zhang et al. 2024).

Rigorous empirical methods (including controlled experiments and human simulations) must be informed by relevant psychological theory (e.g., Gibsonian affordances; see, e.g., Greeno 1994), human factors approaches (e.g., hierarchical task analysis; see, e.g., Stanton 2006), tools (e.g., vigilance protocols, Al-Shargie et al. 2019) and measures (e.g., of situational awareness and cognitive load; Zhang, Yang, Liang, Pitts, Prakah-Asante, Curry, Duerstock, Wachs & Yu 2023, Haapalainen et al. 2010), and tipping point analytics (e.g., autocorrelation and critical slowing down measures, Scheffer et al. 2009). Frameworks for technology acceptance and adoption (e.g., “designing for appropriate resilience and responsivity”, Chiou & Lee 2023) can be employed to measure trust in new technology. Of equal importance is being able to optimally measure loss of trust (which can and will happen, e.g., due a negative experience) and crucially how to restore it—all of which will likely involve ensuring that human-centred design from prototype to deployment considers factors including system accessibility, functionality, usability and adaptability.

Combining the three strands ensures that research in each domain translates into usable, trustworthy demonstrator systems supported by insights into smart agent adoption, trust and trust restoration, and that methods for anticipating and influencing collective change inform interaction design principles for practitioners developing and employing AI systems for collective intelligence across multiple socio-technical domains.

## 5.2 Infrastructure and Governance

A second cross-cutting issue vital to deploying AI for collective intelligence at scale in any use domain is ensuring that such systems have robust infrastructure and governance (I&G) guided by appropriate regulations and principles, delivering trustworthy AI systems and solutions that are human-centred, fair, transparent, and interpretable for the diverse range of end users.

New I&G tools and guidelines for national scale collective AI systems that ensure privacy, quality and integrity of data, and control access to data and systems across relevant AI infrastructures (Aarestrup et al. 2020, Shi et al. 2022, Cao et al. 2023) should be informed by experiences with previous large-scale AI platforms, such as that developed within the EU-wide MIDAS project (Black et al. 2019). One key context in which to pursue this challenge is work developing effective applied AI for national-scale healthcare systems (e.g., cancer, dementia, arthritis; Tedesco et al. 2021, Behera et al. 2021, Henderson et al. 2021). The data sensitivity and outcome criticality of these challenges for AI makes this an ideal domain in which to develop effective I&G tools and thinking. However, considering these issues across multiple diverse applications domains also enables the unique and novel aspects of those settings to inform new thinking on I&G questions.

Such research must address ethical, legal, and social aspects (ELSA) of I&G (Van Veenstra et al. 2021), embedding ELSA accountability in robust governance frameworks and data infrastructure plans. For instance, studies should embed ELSA, FAIR Principles<sup>20</sup> and the Assessment List for Trustworthy Artificial Intelligence (ALTAI)<sup>21</sup> into their Data Management Plans (DMPs) and infrastructure governance and should be informed by outputs of the Artificial Intelligence

<sup>20</sup><https://www.go-fair.org/fair-principles/>

<sup>21</sup><https://digital-strategy.ec.europa.eu/en/library/assessment-list-trustworthy-artificial-intelligence-altai-self-assessment>

Safety Institute (AISI)<sup>22</sup> and other relevant guidelines, e.g., the EU’s Ethics Guidelines for Trustworthy AI,<sup>23</sup> and relevant regulation frameworks such as the EU AI Act.<sup>24</sup>

Finally, the prospect of truly national-scale AI systems of the kinds being considered here foregrounds the pressing need for truly *trans-national* governance structures and mechanisms. These are particularly relevant in collective intelligence settings, since the people, diseases, finance, etc., at the heart of such systems, and the data pertaining to them, all transcend national boundaries. As the Final Report of the United Nation’s AI Advisory Body puts it “the technology is borderless”, necessitating the establishment of “a new social contract for AI that ensures global buy-in for a governance regime that protects and empowers us all” (UN AI Advisory Body 2024).

## 6 Research Strategy

To make significant progress across the research strands outlined above, an effective AI for collective intelligence research strategy must also consider the set of meta-level research challenges that must be overcome if academic research findings are to translate into effective and impactful real-world outcomes. These include addressing underpinning issues around stakeholder engagement; equality, diversity and inclusion (EDI); environmental sustainability; and responsible research and innovation (RRI).

### 6.1 Stakeholder Engagement

We distinguish here between three categories of research stakeholder relevant to AI for collective intelligence research: data partners, skills partners and academic partners. These categories are not disjoint since a single organisation may play more than one of these roles, but they do serve to distinguish between different kinds of research interaction that may be necessary in order to achieve successful applied research in the AI for collective intelligence space at national or trans-national scale.

*Data Partners* are “problem owning” organisations willing to provide controlled access to data, expertise, tools, personnel and strategic guidance relevant to a societal challenge or user need that can be addressed by AI for collective intelligence research. For national-scale efforts, these will tend to be national or trans-national agencies (e.g. the UK’s National Health Service or the UK Health Security Agency) and departments within national government (e.g., the UK’s Department for Health and Social Care), but may also include commercial outfits such as pharmaceutical firms involved in vaccine development, etc. Crucial issues for research collaboration here include those surrounding intellectual property, privacy, regulatory frameworks (e.g., GDPR in the EU), secure data hosting, etc. There are also challenges around the emerging role of synthetic data as a substitute for data that is too sensitive to share or is too hard to anonymize. Such synthetic data can be useful where a mature understanding of the underlying real-world system and the data generating process is in place, but can be problematic in the absence of such an understanding since it can be difficult to provide assurances that the synthetic data captures all of the necessary structural relationships that are present in the original (poorly understood) dataset (Whitney & Norman 2024). More generally, issues around incomplete or noisy data or data that is not sufficiently representative of the underlying population are familiar problems that have significance here.

*Skills Partners* are “problem solving” organisations that are already involved in pioneering the AI and collective intelligence skills, tools and technologies that are driving the next generation of AI for collective intelligence applications, e.g., multi-agent systems, collective systems data science, advanced modelling and machine learning, AI governance and ethics, etc. These may include blue-chip outfits and national facilities (e.g., the UK’s Office of National Statistics)

<sup>22</sup><https://www.gov.uk/government/publications/ai-safety-institute-overview/introducing-the-ai-safety-institute>

<sup>23</sup><https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>

<sup>24</sup><https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206>

but will also include many of the small and medium-sized enterprises (SMEs) emerging in this space (e.g., Flowminder,<sup>25</sup> who leverage decentralised mobility data to support humanitarian interventions in real-time). Research opportunities here include connecting innovating skills partners to the data partners that require their expertise while navigating the intellectual property and commercial sensitivity issues that surround an emerging (and therefore somewhat contested) part of the growing AI consultancy sector.

*Academic Partners* are individuals, research groups or larger academic research activities that are operating in the AI for collective intelligence space. This is a growing area of activity and a key challenge here is connecting and consolidating the emerging community and linking it effectively with the two categories of non-academic stakeholder described above. One key challenge for academic research in this space is balancing the need for rigorous well-understood and mature theory and methods in order to provide quality assurances and guarantee robustness of AI system behaviour against the need to explore and develop new and improved theory and methods that take us beyond the current limited capabilities of extant tools and systems.

### 6.2 Equality, Diversity and Inclusion

The AI workforce lacks diversity (e.g., Young et al. 2021). Moreover, AI technology can tend to impose and perpetuate societal biases (e.g., Kotek et al. 2023). Consequently, it is important that AI for collective intelligence research *operations* be a beacon for best practice in equality, diversity and inclusion (EDI). Moreover, an effective AI for collective intelligence research strategy should itself also be driven by equality, diversity and inclusion *research considerations*. The emerging “AI Divide” separating those that have access to, and command of, powerful new AI technologies from those that do not threatens to further marginalise under-represented, vulnerable, and oppressed individuals and communities (Wang et al. 2024). Collective intelligence methods sometimes focus on achieving consensus and collective agreement (which can tend to prioritise majority views and experiences). However, in addition to aggregating signals at a population level in order to inform the high-level policies and operations of national agencies (which will themselves benefit from being sensitive to population heterogeneity), the AI for collective intelligence research described here is equally interested in deriving bespoke guidance for individuals (or groups) that respects their specific circumstances and needs. This aspect of the research strategy explicitly foregrounds the challenge of reaching and supporting diverse users and those that are intersectionally disadvantaged, e.g., diabetes patients that also have mental health conditions (Benton et al. 2023).

### 6.3 Environmental Sustainability

The carbon footprint of most academic research is dominated by travel (Achten et al. 2013). Consequently, research in this area, like any other, should seek to minimise the use of flights and consider virtual or hybrid meetings wherever possible. Other steps that can be taken to reduce the environmental impact of research practice and move towards “net zero” and “nature positive” ways of working include making sustainable choices for procurement (e.g., accredited sustainable options) and catering (e.g., plant-based food choices).

Moreover, the environment can itself be the focus of AI for collective intelligence research (see “Environmental Intelligence”, above) or a key driving factor (see “Smart City Design”, above). Many associated research themes are consistent with a sustainability agenda in their motivation to achieve effective interventions at scale without consuming vast resources. Intended outcomes and technologies aim to transition society to more sustainable practices (e.g., by using healthcare resource more efficiently, by encouraging sustainable cities, etc.).

However, in common with AI research more generally, AI for collective intelligence makes use of energy-intensive technologies. Computational research is energy-intensive: machine

<sup>25</sup><https://www.flowminder.org/>

learning incurs high CPU/GPU energy cost for training (Patterson et al. 2021), while storage/transfer/duplication of large datasets consumes energy in data centres. Hardware components use rare metals linked to environmental damage and inequalities. Consequently, the environmental impacts of AI for collective intelligence research should always be considered and reduced using, inter alia, computational resources powered by renewable energy, energy efficient algorithms and coding practices, and minimal data duplication. The technologies developed through this kind of research should be evaluated using full Life Cycle Assessment (LCA) techniques that measure their direct impacts (e.g., production, use and disposal costs), indirect impacts (e.g., rebound effects that increase carbon emissions elsewhere in the economy), and identify possible mitigations (e.g., substitution and optimisation effects) (Preist et al. 2019).

#### 6.4 *Responsible Research and Innovation*

Researchers can never know with certainty what future their work will produce, but they can agree on what kind of future they are aiming to bring about, and work inclusively towards making that happen (Owen et al. 2013, Stilgoe et al. 2020). For AI for collective intelligence research, this means working with diverse end users and stakeholders to produce a future in which national-scale AI for collective intelligence systems are tools for societal good (Leonard & Levin 2022).

There are several reasons for taking responsible research and innovation (RRI) concerns especially seriously in the context of artificial intelligence research projects. First, since AI is one of the 17 sensitive research areas named in the UK’s National Security and Investment Act,<sup>26</sup> particular care must be taken by UK universities when establishing and pursuing AI research collaborations. Stakeholder partners must be vetted, and input from the UK Government’s Research Collaboration Advisory Team (RCAT)<sup>27</sup> must be sought where there are concerns regarding, for instance, the exploitation of intellectual property arising from the research activity. Moreover, applied AI research is often fuelled by data that is sensitive, meaning that huge care must be taken to safeguard this data and ensure privacy through the use of, e.g., secure research data repositories that feature robust controlled data access protocols. More generally, since AI innovations have the potential to radically reshape the future in ways that are very hard to predict, articulating a clear shared vision of the future that is being aimed for is particularly important. Finally, AI researchers have a responsibility to engage with the public discourse around AI which is currently driving considerable anxiety and confusion.<sup>28</sup>

## 7 Unifying Research Challenges

The research strategy outlined here sets out to develop, build and evaluate systems that exploit machine learning and AI to achieve improved collective intelligence at multiple scales: driving improved policy and operations at the level of national agencies and offering bespoke guidance and decision support to individual citizens.

Why are systems of this kind not already in routine operation? Many of the component technologies are established and some are becoming reasonably well understood: recommender systems (Resnick & Varian 1997), machine learning at scale (Lwakatare et al. 2020), networked infrastructure and Internet devices (Radanliev et al. 2020, Rashid et al. 2023), conversational AI (Kulkarni et al. 2019), network science analyses (Börner et al. 2007), etc. However, several important and interacting challenges obstruct the realisation of AI for collective intelligence and these must be targets for this research effort. Here, we distinguish three categories: *human* challenges, *technical* challenges and *scale* challenges. Each application domain manifests a

<sup>26</sup><https://www.gov.uk/government/publications/national-security-and-investment-act-guidance-on-notifiable-acquisitions>

<sup>27</sup><https://www.gov.uk/government/organisations/research-collaboration-advice-team>

<sup>28</sup>[https://www.ey.com/en\\_us/consulting/businesses-can-stop-rising-ai-use-from-fueling-anxiety](https://www.ey.com/en_us/consulting/businesses-can-stop-rising-ai-use-from-fueling-anxiety)

**Table 1** Examples of how three different categories of unifying research challenge apply within five different AI for collective intelligence application domains.

Application Domain	Human Challenges	Technical Challenges	Scale Challenges
Smart City Design	Delivering informative visualization of complex city dynamics	Predicting outcomes of interventions that trigger behavioural change	Balancing short-term optimisation and medium term sustainability
Pandemic Resilience	Achieving and maintaining trust in guidance/interventions	Predicting reproductive number under phased vaccination efforts	Differentiating regional and national collective intelligence
Environmental Intelligence	Providing transparent provenance indications for complex climate data	Evaluating climate avatar performance in the wild	Coping with bursty demand during extreme weather events
Financial Stability	Incorporating user values in personal finance recommendations	Achieving privacy-preserving ML for financial analytics	Operating at speeds consistent with flash crash cascade dynamics
Healthcare Ecosystems	Achieving effective usability for hard to reach patients	Identifying continuously evolving clusters of similar patient types	Achieving persistent uptake at scale for patients and NHS trusts

combination of challenges in a distinctive way (see Table 1), but since these challenges are inter-related they should be approached holistically.<sup>29</sup>

### 7.1 Human Challenges

In order to be successful, the systems developed at the boundary between AI and collective intelligence must actually be engaged with and used by individual people as well as by institutions and agencies. For this to be the case, these systems must be trusted. Individual people must trust the systems with their data and must trust the guidance that they are offered.<sup>30</sup> Institutions and agencies must trust the reliability of the aggregated findings delivered by the systems and must trust that the systems will operate in a way that does not expose them to reputational risk by disadvantaging users or putting them at risk. The European Commission’s High-Level Expert Group on Artificial Intelligence suggests that trust in AI should arise from seven properties: empowering human agency, security, privacy, transparency, fairness, value alignment, and accountability.<sup>31</sup>

What are the hallmarks of these properties that users of AI for collective intelligence systems intuitively and readily recognise? What kinds of guarantees for these properties could be credibly offered to regulators or law makers? More generally, how can users of all kinds become confident that these properties are present in a particular system and remain confident as they continue to interact with it?

Moreover, since the systems envisioned here purport to offer *bespoke* decision support tailored to the needs of individual users, one acute aspect of this challenge relates to supporting the needs

<sup>29</sup>One theoretical framework with promising potential to support and inter-relate the challenges being considered here is that offered by studies in cumulative cultural evolution (Mesoudi & Thornton 2018, Smaldino 2014).

<sup>30</sup>There is some debate as to whether notions of ‘trust’ and ‘trustworthiness’ are appropriate for framing the legitimacy of AI systems; compare, for example, the work of Andras et al. (2018) with the position of Bryson (2018). Even the use of a term like ‘guidance’ to describe the kind of support that AI systems might be designed to offer can be reminiscent of previous attempts to shift public behaviour through the ‘libertarian paternalism’ of nudge economics, an approach that was discredited precisely because of its tendency to disempower or even coerce people rather than fully inform or partner with them (Goodwin 2012).

<sup>31</sup><https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>

of all kinds of user including those from marginalised or under-represented groups. It is typically the case that machine learning extracts patterns that generalise over the diversity in a data set in order to capture central tendencies, robust trends, etc. This can fail to capture, respect or represent the features of dataset outliers. Within society, these outliers are often individuals particularly in need of support, and this support may not be useful unless it is sensitive to the specific features of these individuals' circumstances. Meeting this challenge, and the more general challenge of deserving, achieving and maintaining trust, will require an interdisciplinary combination of both social and technical insights (Gilbert & Bullock 2014).

## 7.2 *Technical Challenges*

Amongst the many technical challenges that must be overcome to enable AI for collective intelligence to be effective, we will mention three: nonstationarity in collective systems, privacy and robustness of multi-level machine learning, and the ethics of multi-agent collective decision support.

All machine learning makes a gamble that the future will resemble the past, yet we know that the data from collective systems can be nonstationary, i.e., these systems can make transitions between regimes that may differ radically from one another (Scheffer et al. 2009). How can we anticipate and detect these sudden shifts, phase transitions, regime changes, and tipping points at the level of entire collectives, sub-groups and even individuals? These questions have been considered within collective intelligence research (Mann 2022, Tilman et al. 2023), and the large scale of systems under study here offers potential to trial methods that have been applied to physical and biological systems, e.g., early warning signals from dynamical systems theory (Scheffer et al. 2009), but can these be effective for complex fast-moving socio-technical systems?

Amongst the many other machine learning challenges relevant here, we will highlight two that arise as a consequence of the fundamentally multi-level nature of collective intelligence. The AI4CI Loop (Fig. 1) depicts the way in which the approach to AI for collective intelligence being pursued here involves machine learning models that deliver findings at different levels of description, from findings that characterise the entire collective through intermediate results related to sub-groups within the collective to bespoke results relevant to individual members of the collective. Delivering this requires (likely unsupervised) methods to cluster and/or unbundle heterogeneous data stream trajectories derived from groups and individuals. Moreover, achieving this whilst maintaining the privacy of individual members of the collective requires robust privacy-preserving machine learning methods. Employing foundation models to capture and compress the patterns in the collective system is one possible approach, but understanding the vulnerabilities of these models remains an open research challenge (Chen et al. 2023, Messeri & Crockett 2024).

Within collectives, one member's actions can affect other members. In this context, systems that support decision making do not only impact their direct user (Ajmeri et al. 2018, Vinitzky et al. 2023). While there is potential to leverage this collective decision making to achieve efficient coordinated outcomes (Jacyno et al. 2009), it remains the case that a typical AI agent tends to cater to the interests of their primary user even if they are intended to reflect the preferences of multiple stakeholders (Murukannaiah et al. 2020). This may reinforce existing privileges and could worsen the challenges faced by vulnerable individuals and marginalised groups. Thus, it is imperative that AI agents consider and communicate the broader collective implications of the decision support that they offer. In this way we can encourage these agents to respect societal norms and their stakeholders' needs and value preferences, and inform decisions that promote fairness, inclusivity, sustainability and equitability (Murukannaiah et al. 2020, Woodgate & Ajmeri 2022, 2024).

### 7.3 Scale Challenges

This paper has articulated a set of research challenges in terms of “producing national-scale AI for collective intelligence”. For some domains, e.g., pandemic response, this scale might appear to be a natural level of description because relevant policy, operations, data and governance are all ultimately defined at the level of national government. However, most if not all collective intelligence challenges engage with multiple spatial, social and governmental scales. Decision making at national, regional, local, household and personal scales are simultaneously in play, and in some cases trans-national scales are also significant as when pandemics, environmental disasters or financial contagion cross national borders. Consequently, the adjective “national-scale” should not be taken here to imply a single scale of operation and a single locus of decision making. Rather, the most effective AI for collective intelligence systems will be able to operate in a hierarchical, cross-scale fashion as anticipated in the work of, e.g., Ostrom (2010) and others.

Operating at national or trans-national scale can help to address some of the human and technical challenges discussed in the previous sections: national agencies, such as the UK’s NHS or Met Office, are often trusted agencies; engaging with large populations of users can enable better support for marginalised groups of users that are typically under-represented; operating at scale can increase sensitivity to nonstationarity in an underlying collective system. However, scale also brings its own challenges in terms of establishing and managing appropriate infrastructure in a way that is secure, well-governed and sustainable.

Such infrastructure must support data collection at massive volume. Services must be delivered at point of use, in real time, without failure. Sensitive personal data must be handled securely and systems must respect the privacy of individuals while also providing solutions that rely on data aggregation and sharing. Since the security of national digital infrastructure is increasingly important in a global context where cyberattacks and information operations are becoming more common, infrastructure and service delivery must be robust to external threats as well as internal perturbations and flaws. The infrastructure and associated services and processes must be subject to appropriate and effective governance. Finally, environmental sustainability must be a core aim. How can the national scale systems and services envisioned here operate in a way that has low environmental impact?

## 8 Conclusion

There is considerable potential for productive research at the intersection between the fields of artificial intelligence and collective intelligence. This paper has presented one research strategy for operating at this intersection, articulated in terms of the AI4CI Hub’s research vision and research themes, its approach to prosecuting interdisciplinary, collaborative research, and its set of unifying research challenges. The Hub’s first steps include pursuing case study research projects within each of the AI4CI themes in collaboration with relevant stakeholders,<sup>32</sup> hosting workshops and symposia to cross-fertilise and disseminate new tools, methods and thinking at the AI/collective intelligence interface,<sup>33</sup> and launching a funding opportunity to support new collaborative research activities in this space across the UK.<sup>34</sup> No doubt there are many viable alternative research strategies at this same interface, and we echo Nesta’s conclusion that “the field can only evolve through more organisations experimenting with different models of AI and CI and the opportunity to deliver novel solutions to real-world challenges”.<sup>35</sup>

<sup>32</sup><https://ai4ci.ac.uk/research/>

<sup>33</sup><https://ai4ci.ac.uk/events/>

<sup>34</sup><https://ai4ci.ac.uk/funding/>

<sup>35</sup><https://www.nesta.org.uk/report/future-minds-and-machines/>

### Acknowledgements

This work was supported by UKRI EPSRC Grant No. EP/Y028392/1: *AI for Collective Intelligence (AI4CI)*.

### Competing Interests

The authors declare none.

### References

- Aarestrup, F., Albeyatti, A., Armitage, W., Auffray, C., Augello, L., Balling, R., Benhabiles, N., Bertolini, G., Bjaalie, J., Black, M., Blomberg, N., Bogaert, P., Bubak, M., Claerhout, B., Clarke, L., De Meulder, B., D’Errico, G., Di Meglio, A., Forgo, N., Gans-Combe, C., Gray, A., Gut, I., Gyllenberg, A., Hemmrich-Stanisak, G., Hjorth, L., Ioannidis, Y., Jarmalaite, S., Kel, A., Kherif, F., Korbel, J., Larue, C., László, M., Maas, A., Magalhaes, L., Manneh-Vangramberen, I., Morley-Fletcher, E., Ohmann, C., Oksvold, P., Oxtoby, N., Perseil, I., Pezoulas, V., Riess, O., Riper, H., Roca, J., Rosenstiel, P., Sabatier, P., Sanz, F., Tayeb, M., Thomassen, G., Van Bussel, J., Van Den Bulcke, M. & Van Oyen, H. (2020), ‘Towards a European Health Research and Innovation Cloud (HRIC)’, *Genome Medicine* **12**, 1–14. DOI:<https://dx.doi.org/10.1186/s13073-020-0713-z>.
- Acar, O. A. (2023), ‘Crowd science and science skepticism’, *Collective Intelligence* **2**(1). DOI:<https://dx.doi.org/10.1177/263391372311764>.
- Achten, W. M. J., Almeida, J. & Muys, B. (2013), ‘Carbon footprint of science: More than flying’, *Ecological Indicators* **34**, 352–355. DOI:<https://dx.doi.org/10.1016/j.ecolind.2013.05.025>.
- Ajmeri, N., Guo, H., Murukannaiah, P. K. & Singh, M. P. (2018), 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)’, IJCAI, Stockholm, pp. 28–34. DOI:<https://dx.doi.org/10.24963/ijcai.2018/4>.
- Al-Shargie, F., Tariq, U., Mir, H., Alawar, H., Babiloni, F. & Al-Nashash, H. (2019), ‘Vigilance decrement and enhancement techniques: A review’, *Brain Sciences* **9**(8), 178–203. DOI:<https://dx.doi.org/10.3390/brainsci9080178>.
- An, L., Grimm, V., Bai, Y., Sullivan, A., Turner II, B., Malleon, N., Heppenstall, A., Vincenot, C., Robinson, D., Ye, X., Liu, J., Lindkvist, E. & Tang, W. (2023), ‘Modeling agent decision and behavior in the light of data science and artificial intelligence’, *Environmental Modelling & Software* **166**, 105713. DOI:<https://dx.doi.org/10.1016/j.envsoft.2023.105713>.
- Andras, P., Esterle, L., Guckert, M., Han, T. A., Lewis, P. R., Milanovic, K., Payne, T., Perret, C., Pitt, J., Powers, S. T., Urquhart, N. & Wells, S. (2018), ‘Trusting intelligent machines: Deepening trust within socio-technical systems’, *IEEE Technology and Society Magazine* **37**(4), 76–83. DOI:<https://dx.doi.org/10.1109/MTS.2018.2876107>.
- Arner, D. W., Barberis, J. & Buckley, R. P. (2015), The evolution of FinTech: A new post-crisis paradigm?, Technical Report 2015/047, Hong Kong: University of Hong Kong, Faculty of Law. DOI:<https://dx.doi.org/10.2139/ssrn.2676553>.
- Batty, M. (2024), ‘AI and design’, *Environment and Planning B: Urban Analytics and City Science* **51**, 23998083241236619. DOI:<https://dx.doi.org/10.1177/23998083241236619>.
- Batty, M., Clifton, J., Tyler, P. & Wan, L. (2020), ‘The post-Covid city’, *Cambridge Journal of Regions, Economy and Society* **15**(3), 447–457. DOI:<https://dx.doi.org/10.1093/cjres/rsac041>.



- Batty, M., Crooks, A. T., See, L. M. & Heppenstall, A. J. (2012), Perspectives on agent-based models and geographical systems, in A. J. Heppenstall, A. T. Crooks, L. M. See & M. Batty, eds, 'Agent-Based Models of Geographical Systems', Springer, Dordrecht, pp. 1–15. DOI:<https://dx.doi.org/10.1007/978-90-481-8927-4>.
- Baxter, G. & Sommerville, I. (2011), 'Socio-technical systems: From design methods to systems engineering', *Interacting with Computers* **23**(1), 4–17. DOI:<https://doi.org/10.1016/j.intcom.2010.07.003>.
- Bazarbash, M. (2019), FinTech in financial inclusion: Machine learning applications in assessing credit risk, Technical Report 2019/109, International Monetary Fund. DOI:<https://dx.doi.org/10.5089/9781498314428.001>.
- Behera, C., Condell, J., Dora, S., Gibson, D. & Leavey, G. (2021), 'State-of-the-art sensors for remote care of people with dementia during a pandemic: A systematic review', *Sensors* **21**(14). DOI:<https://dx.doi.org/10.3390/s21144688>.
- Benton, M., Cleal, B., Prina, M., Baykoca, J., Willaing, I., Price, H. & Ismail, K. (2023), 'Prevalence of mental disorders in people living with type 1 diabetes: A systematic literature review and meta-analysis', *General Hospital Psychiatry* **80**, 1–16. DOI:<https://dx.doi.org/10.1016/j.genhosppsych.2022.11.004>.
- Berditchevskaia, A. & Baeck, P. (2020), 'The future of minds and machines: How artificial intelligence can enhance collective intelligence', NESTA: London, UK. <https://www.nesta.org.uk/report/future-minds-and-machines>.
- Berditchevskaia, A., Maliaraki, E. & Stathoulopoulos, K. (2022), 'A descriptive analysis of collective intelligence publications since 2000, and the emerging influence of artificial intelligence', *Collective Intelligence* **1**(1). DOI:<https://dx.doi.org/10.1177/26339137221107924>.
- Black, M., Wallace, J., Rankin, D., Carlin, P., Bond, R., Mulvenna, M., Cleland, B., Fischaber, S., Epelde, G., Nikolic, G., Pajula, J. & Connolly, R. (2019), Meaningful integration of data, analytics and services of computer-based medical systems: The MIDAS touch, in 'Proceedings of the 32nd IEEE International Symposium on Computer-Based Medical Systems (CBMS)', pp. 104–105. DOI:<https://dx.doi.org/10.1109/CBMS.2019.00031>.
- Bonabeau, E., Dorigo, M. & Theraulaz, G. (1999), *Swarm Intelligence: From Natural to Artificial Systems*, Oxford University Press. DOI:<https://doi.org/10.1093/oso/9780195131581.001.0001>.
- Börner, K., Sanyal, S. & Vespignani, A. (2007), Network science, in B. Cronin, ed., 'Annual Review of Information Science & Technology', Information Today Inc./American Society for Information Science and Technology, Medford, NJ, chapter 12, pp. 537–607.
- Bourazeri, A. & Pitt, J. (2018), 'Collective attention and active consumer participation in community energy systems', *International Journal of Human-Computer Studies* **119**. DOI:<https://doi.org/10.1016/j.ijhcs.2018.06.001>.
- Bratteteig, T., Bødker, K., Dittrich, Y., Mogensen, P. H. & Simonsen, J. (2012), Methods: Organising principles and general guidelines for participatory design projects, in J. Simonsen & T. Robertson, eds, 'Routledge International Handbook of Participatory Design', Routledge, London and New York, pp. 117–144.
- Brooks-Pollock, E., Danon, L., Jombart, T. & Pellis, L. (2021), 'Modelling that shaped the early COVID-19 pandemic response in the UK', *Philosophical Transactions of the Royal Society of London, Series B* **376**(1829), 3762021000120210001. DOI:<https://dx.doi.org/10.1098/rstb.2021.0001>.

- Bryson, J. (2018), AI & global governance: No one should trust AI, Blog Post, United Nations University Centre for Policy Research.
- Buckle, M., Chen, J., Guo, Q. & Li, X. (2023), ‘Does smile help detect the UK’s price leadership change after MiFID?’, *International Review of Economics & Finance* **84**, 756–769. DOI:<https://dx.doi.org/10.1016/j.iref.2022.11.033>.
- Bullock, S. & Sayama, H. (2023), Agent heterogeneity mediates extremism in an adaptive social network model, in H. Iizuka, K. Suzuki, R. Uno, L. Damiano, N. Sychalav, M. Aguilera, E. Izquierdo, R. Suzuki & M. Baltieri, eds, ‘Proceedings of the Artificial Life Conference 2023 (ALIFE 2023)’, MIT Press. DOI:[https://doi.org/10.1162/isal\\_a\\_00628](https://doi.org/10.1162/isal_a_00628).
- Cao, H., Wachowicz, M., Richard, R. & Hsu, C.-H. (2023), ‘Fostering new vertical and horizontal IoT applications with intelligence everywhere’, *Collective Intelligence* **2**(4). DOI:<https://dx.doi.org/10.1177/26339137231208966>.
- Cartlidge, J., Szostek, C., Luca, M. D. & Cliff, D. (2012), Too fast too furious: Faster financial-market trading agents can give less efficient markets, in J. Filipe & A. L. N. Fred, eds, ‘Proceedings of 4th International Conference on Agents and Artificial Intelligence (ICAART)’, SciTePress, Setúbal, Portugal, pp. 126–135. DOI:<https://dx.doi.org/10.5220/0003720301260135>.
- Challen, R., Denny, J., Pitt, M., Gompels, L., Edwards, T. & Tsaneva-Atanasova, K. (2019), ‘Artificial intelligence, bias and clinical safety’, *BMJ Quality & Safety* **28**(3), 231–237. DOI:<https://dx.doi.org/10.1136/bmjqs-2018-008370>.
- Challen, R., Tsaneva-Atanasova, K., Pitt, M., Edwards, T., Gompels, L., Lacasa, L., Brooks-Pollock, E. & Danon, L. (2021), ‘Estimates of regional infectivity of COVID-19 in the United Kingdom following imposition of social distancing measures’, *Philosophical Transactions of the Royal Society of London, Series B* **376**(1829), 3762020028020200280. DOI:<https://dx.doi.org/10.1098/rstb.2020.0280>.
- Chen, C. & Campbell, N. D. F. (2022), Analysing training-data leakage from gradients through linear systems and gradient matching, in ‘The 33rd British Machine Vision Conference (BMVC 2022)’, BMVA Press.
- Chen, C., Namboodiri, V. P. & Padget, J. (2023), Understanding the vulnerability of CLIP to image compression, in ‘Proceedings of the Workshop on Robustness of Few-shot and Zero-shot Learning in Foundation Models (NeurIPS 2023)’. <https://arxiv.org/abs/2311.14029>.
- Chiou, E. K. & Lee, J. D. (2023), ‘Trusting automation: Designing for responsivity and resilience’, *Human Factors* **65**(1), 137–165. DOI:<https://dx.doi.org/10.1177/001872082111009>.
- Choung, H., David, P. & Ross, A. (2023), ‘Trust in AI and its role in the acceptance of AI technologies’, *International Journal of Human-Computer Interaction* **39**(9), 1727–1739. DOI:<https://dx.doi.org/10.1080/10447318.2022.2050543>.
- Dambanemuya, H. K., Wachs, J. & Ágnes Horvát, E. (2023), Understanding (ir)rational herding online, in M. Bernstein, S. Savage & A. Bozzon, eds, ‘Proceedings of The ACM Collective Intelligence Conference (CI)’, ACM, Delft, pp. 79–88. DOI:<https://dx.doi.org/10.1145/3582269.3615598>.
- Duckworth, C., Guy, M. J., Kumaran, A., O’Kane, A. A., Ayobi, A., Chapman, A., Marshall, P. & Boniface, M. (2024), ‘Explainable machine learning for real-time hypoglycemia and hyperglycemia prediction and personalized control recommendations’, *Journal of Diabetes Science and Technology* **18**(1), 113–123. DOI:<https://dx.doi.org/10.1177/19322968221103561>.

- Emanuel, E. J. & Wachter, R. M. (2019), ‘Artificial intelligence in health care: Will the value match the hype?’, *JAMA* **321**(23), 2281–2282. DOI:<https://dx.doi.org/10.1001/jama.2019.4914>.
- Gilbert, N. & Bullock, S. (2014), ‘Complexity at the social science interface’, *Complexity* **19**(6), 1–4. DOI:<https://dx.doi.org/10.1002/cplx.21550>.
- Góis, A. R., Santos, F. P., Pacheco, J. M. & Santos, F. C. (2019), ‘Reward and punishment in climate change dilemmas’, *Scientific Reports* **9**(1), 16193. DOI:<https://doi.org/10.1038/s41598-019-52524-8>.
- Goodwin, T. (2012), ‘Why we should reject “nudge”’, *Politics* **32**(2), 85–92. DOI:<https://doi.org/10.1111/j.1467-9256.2012.01430.x>.
- Greeno, J. G. (1994), ‘Gibson’s affordances’, *Psychological Review* **101**(2), 336–342. DOI:<https://dx.doi.org/10.1037/0033-295X.101.2.336>.
- Haapalainen, E., Kim, S., Forlizzi, J. F. & Dey, A. K. (2010), Psycho-physiological measures for assessing cognitive load, in ‘Proceedings of the 12th ACM International Conference on Ubiquitous Computing (UbiComp ’10)’, Association for Computing Machinery, New York, NY, USA, p. 301–310. DOI:<https://dx.doi.org/10.1145/1864349.1864395>.
- Hart, S., Banks, V., Bullock, S. & Noyes, J. (2022), Understanding human decision-making when controlling UAVs in a search and rescue application, in T. Ahram & R. Taiar, eds, ‘Human Interaction & Emerging Technologies (IHJET 2022): Artificial Intelligence & Future Applications. AHFE (2022) International Conference’, AHFE Open Access, vol 68. AHFE International, USA. DOI:<http://doi.org/10.54941/ahfe1002768>.
- Henderson, J., Condell, J., Connolly, J., Kelly, D. & Curran, K. (2021), ‘Review of wearable sensor-based health monitoring glove devices for rheumatoid arthritis’, *Sensors* **21**(5), 1–32. DOI:<https://dx.doi.org/10.3390/s21051576>.
- Jacyno, M., Bullock, S., Luck, M. & Payne, T. R. (2009), Emergent service provisioning and demand estimation through self-organizing agent communities, in C. Sierra, C. Castelfranchi, K. S. Decker & J. S. Sichman, eds, ‘Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)’, ACM, Budapest, pp. 481–488. DOI:<https://dx.doi.org/10.1145/1558013.1558079>.
- Johnson, N., Zhao, G., Hunsader, E., Qi, H., Johnson, N., Meng, J. & Tivnan, B. (2013), ‘Abrupt rise of new machine ecology beyond human response time’, *Scientific Reports* **3**(1), 2627. DOI:<https://dx.doi.org/10.1038/srep02627>.
- Klein, T. (2022), ‘A note on GameStop, short squeezes, and autodidactic herding: An evolution in financial literacy?’, *Finance Research Letters* **46**, 102229. DOI:<https://dx.doi.org/10.1016/j.frl.2021.102229>.
- Koldunov, N. & Jung, T. (2024), ‘Local climate services for all, courtesy of large language models’, *Communications Earth & Environment* **5**(1), 13. DOI:<https://dx.doi.org/10.1038/s43247-023-01199-1>.
- Kotek, H., Dockum, R. & Sun, D. Q. (2023), Gender bias and stereotypes in large language models, in M. Bernstein, S. Savage & A. Bozzon, eds, ‘Proceedings of The ACM Collective Intelligence Conference (CI)’, ACM, Delft, pp. 12–24. DOI:<https://dx.doi.org/10.1145/3582269.3615599>.
- Kulkarni, P., Mahabaleshwar, A., Kulkarni, M., Sirsakar, N. & Gadgil, K. (2019), Conversational AI: An overview of methodologies, applications & future scope, in ‘Proceedings of the 5th International Conference On Computing, Communication, Control And Automation (ICCUBEA)’, IEEE, Pune, pp. 1–7. DOI:<https://dx.doi.org/10.1109/ICCUBEA47591.2019>.

- Lee, J. D. & See, K. A. (2004), ‘Trust in automation: Designing for appropriate reliance’, *Human Factors* **46**(1), 50–80. DOI:[https://dx.doi.org/10.1518/hfes.46.1.50\\_30392](https://dx.doi.org/10.1518/hfes.46.1.50_30392).
- Leonard, N. E. & Levin, S. A. (2022), ‘Collective intelligence as a public good’, *Collective Intelligence* **1**(1). DOI:<https://dx.doi.org/10.1177/26339137221083293>.
- Liu, A., Jahanshahloo, H., Chen, J. & Eshraghi, A. (2023), ‘Trading patterns in the bitcoin market’, *The European Journal of Finance* . DOI:<https://dx.doi.org/10.1080/1351847X.2023.2241883>.
- Lwakatare, L. E., Raj, A., Crnkovic, I., Bosch, J. & Olsson, H. H. (2020), ‘Large-scale machine learning systems in real-world industrial settings: A review of challenges and solutions’, *Information and Software Technology* **127**, 106368. DOI:<https://dx.doi.org/10.1016/j.infsof.2020.106368>.
- Malleson, N., Birkin, M., Birks, D., Ge, J., Heppenstall, A., Manley, E., McCulloch, J. & Ternes, P. (2022), ‘Agent-based modelling for urban analytics: State of the art and challenges’, *AI Communications* **35**(4), 393–406. DOI:<https://dx.doi.org/10.3233/AIC-220114>.
- Mann, R. P. (2022), ‘Collective decision-making under changing social environments among agents adapted to sparse connectivity’, *Collective Intelligence* **1**(2). DOI:<https://dx.doi.org/10.1177/26339137221121347>.
- Mesoudi, A. & Thornton, A. (2018), ‘What is cumulative cultural evolution?’, *Proceedings of the Royal Society of London, Series B* **285**(1880), 20180712. DOI:<https://doi.org/10.1098/rspb.2018.0712>.
- Messeri, L. & Crockett, M. J. (2024), ‘Artificial intelligence and illusions of understanding in scientific research’, *Nature* **627**(8002), 49–58. DOI:<https://dx.doi.org/10.1038/s41586-024-07146-0>.
- Murukannaiah, P. K., Ajmeri, N., Jonker, C. M. & Singh, M. P. (2020), New foundations of ethical multiagent systems, in ‘Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)’, IFAAMAS, Auckland, pp. 1706–1710. DOI:<https://dx.doi.org/10.5555/3398761.3398958>.
- NHS (2019), ‘The NHS long term plan’, NHS. <https://www.longtermplan.nhs.uk/publication/nhs-long-term-plan/>.
- Ostrom, E. (2010), ‘Beyond markets and states: Polycentric governance of complex economic systems’, *Transnational Corporations Review* **2**(2), 1–12. DOI:<https://doi.org/10.1080/19186444.2010.11658229>.
- Owen, R., Stilgoe, J., Macnaghten, P., Gorman, M., Fisher, E., Guston, D. & Bessant, J. (2013), A framework for responsible innovation, in ‘Responsible Innovation: Managing the Responsible Emergence of Science and Innovation in Society’, Wiley, Hoboken, NJ, pp. 27–50. DOI:<https://dx.doi.org/10.1002/9781118551424.ch2>.
- Patterson, D. A., Gonzalez, J., Le, Q. V., Liang, C., Munguia, L., Rothchild, D., So, D. R., Texier, M. & Dean, J. (2021), ‘Carbon emissions and large neural network training’, Pre-print. arXiv:2104.10350. DOI:<https://dx.doi.org/10.48550/arXiv.2104.10350>.
- Pitonakova, L., Crowder, R. & Bullock, S. (2018), ‘The Information-Cost-Reward framework for understanding robot swarm foraging’, *Swarm Intelligence* **12**(1), 71–96. DOI:[10.1007/s11721-017-0148-3](https://doi.org/10.1007/s11721-017-0148-3).

- Preist, C., Schien, D. & Shabajee, P. (2019), Evaluating sustainable interaction design of digital services: The case of YouTube, in 'Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems', ACM, Glasgow, pp. 1–12. DOI:<https://dx.doi.org/10.1145/3290605.3300627>.
- Radanliev, P., De Roure, D., Walton, R., Van Kleek, M., Montalvo, R. M., Santos, O., Maddox, L. & Cannady, S. (2020), 'COVID-19 what have we learned? The rise of social machines and connected devices in pandemic management following the concepts of predictive, preventive and personalized medicine', *EPMA Journal* **11**(3), 311–332. DOI:<https://dx.doi.org/10.1007/s13167-020-00218-x>.
- Rajkomar, A., Dean, J. & Kohane, I. (2019), 'Machine learning in medicine', *New England Journal of Medicine* **380**(14), 1347–1358. DOI:<https://dx.doi.org/10.1056/NEJMra1814259>.
- Rashid, M. T., Wei, N. & Wang, D. (2023), 'A survey on social-physical sensing: An emerging sensing paradigm that explores the collective intelligence of humans and machines', *Collective Intelligence* **2**(2). DOI:<https://dx.doi.org/10.1177/26339137231170825>.
- Resnick, P. & Varian, H. R. (1997), 'Recommender systems', *Communications of the ACM* **40**(3), 56–58. DOI:<https://dx.doi.org/10.1145/245108.245121>.
- Scheffer, M., Bascompte, J., Brock, W. A., Brovkin, V., Carpenter, S. R., Dakos, V., Held, H., van Nes, E. H., Rietkerk, M. & Sugihara, G. (2009), 'Early-warning signals for critical transitions', *Nature* **461**, 53–59. DOI:<https://dx.doi.org/10.1038/nature08227>.
- Shi, X., Nikolic, G., Fischaber, S., Black, M., Rankin, D., Epelde, G., Beristain, A., Alvarez, R., Arrue, M., Pita Costa, J., Grobelnik, M., Stopar, L., Pajula, J., Umer, A., Poliwoda, P., Wallace, J., Carlin, P., Pääkkönen, J. & De Moor, B. (2022), 'System architecture of a European platform for health policy decision making: MIDAS', *Frontiers in Public Health* **10**, 1–13. DOI:<https://dx.doi.org/10.3389/fpubh.2022.838438>.
- Shi, Z. & Cartlidge, J. (2022), State dependent parallel neural Hawkes process for limit order book event stream prediction and simulation, in 'Proceedings of the 28th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)', Washington DC, pp. 1607–1615. DOI:<https://dx.doi.org/10.1145/3534678.3539462>.
- Shi, Z. & Cartlidge, J. (2024), 'Neural stochastic agent-based limit order book simulation with neural point process and diffusion probabilistic model', *Intelligent Systems in Accounting, Finance and Management*. DOI:<https://dx.doi.org/10.1002/isaf.1553>.
- Smaldino, P. E. (2014), 'The cultural evolution of emergent group-level traits', *Behavioral and Brain Sciences* **37**, 243–295. DOI:<https://doi.org/10.1017/S0140525X13001544>.
- Smaldino, P. E. & O'Connor, C. (2022), 'Interdisciplinarity can aid the spread of better methods between scientific communities', *Collective Intelligence* **1**(2). DOI:<https://dx.doi.org/10.1177/26339137221131816>.
- Smith, L. G. E., Blackwood, L. & Thomas, E. F. (2020), 'The need to refocus on the group as the site of radicalization', *Perspectives on Psychological Science* **15**(2), 327–352. DOI:<https://doi.org/10.1177/1745691619885870>.
- Spooner, F., Abrams, J. F., Morrissey, K., Shaddick, G., Batty, M., Milton, R., Dennett, A., Lomax, N., Malleson, N., Nelissen, N., Coleman, A., Nur, J., Jin, Y., Greig, R., Shenton, C. & Birkin, M. (2021), 'A dynamic microsimulation model for epidemics', *Social Science & Medicine* **291**, 114461. DOI:<https://dx.doi.org/10.1016/j.socscimed.2021.114461>.

- Stanton, N. A. (2006), ‘Hierarchical task analysis: Developments, applications, and extensions’, *Applied Ergonomics* **37**(1), 55–79. DOI:<https://dx.doi.org/10.1016/j.apergo.2005.06.003>.
- Stawarz, K., Katz, D., Ayobi, A., Marshall, P., Yamagata, T., Santos-Rodriguez, R., Flach, P. & O’Kane, A. A. (2023), ‘Co-designing opportunities for human-centred machine learning in supporting type 1 diabetes decision-making’, *International Journal of Human-Computer Studies* **173**, 103003. DOI:<https://dx.doi.org/10.1016/j.ijhcs.2023.103003>.
- Stilgoe, J., Owen, R. & Macnaghten, P. (2020), Developing a framework for responsible innovation, in J. S. Andrew Maynard, ed., ‘The Ethics of Nanotechnology, Geoengineering and Clean Energy’, Routledge, London, pp. 347–359. DOI:<https://dx.doi.org/10.4324/9781003075028-22>.
- Tedesco, S., Andrulli, M., Larsson, M. A., Kelly, D., Timmons, S., Alamäki, A., Barton, J., Condell, J., O’Flynn, B. & Nordström, A. (2021), Investigation of the analysis of wearable data for cancer-specific mortality prediction in older adults, in ‘Proceedings of the 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)’, Virtual, pp. 1848–1851. DOI:<https://dx.doi.org/10.1109/EMBC46164.2021.9630370>.
- Tilman, A. R., Vasconcelos, V. V., Akçay, E. & Plotkin, J. B. (2023), ‘The evolution of forecasting for decision-making in dynamic environments’, *Collective Intelligence* **2**(4). DOI:<https://dx.doi.org/10.1177/26339137231221726>.
- Topol, E. (2019), The Topol Review: Preparing the healthcare workforce to deliver the digital future, Technical report, An independent report on behalf of the UK Government’s Secretary of State for Health and Social Care. <https://topol.hee.nhs.uk/wp-content/uploads/HEE-Topol-Review-2019.pdf>.
- Treen, K. M. d., Williams, H. T. P. & O’Neill, S. J. (2020), ‘Online misinformation about climate change’, *Wiley Interdisciplinary Reviews: Climate Change* **11**(5), e665. DOI:<https://dx.doi.org/10.1002/wcc.665>.
- UN AI Advisory Body (2024), Governing AI for Humanity, Final Report, United Nations.
- Vaghefi, S. A., Stambach, D., Muccione, V., Bingler, J., Ni, J., Kraus, M., Allen, S., Colesanti-Senni, C., Wekhof, T., Schimanski, T. et al. (2023), ‘ChatClimate: Grounding conversational AI in climate science’, *Communications Earth & Environment* **4**(1), 480. DOI:<https://dx.doi.org/10.1038/s43247-023-01084-x>.
- van Thiel, D. & Elliott, K. (2024), ‘Responsible access to credit for sole-traders and micro-organisations under unstable market conditions with psychometrics’, *The European Journal of Finance* . Forthcoming. DOI:<https://dx.doi.org/10.1080/1351847X.2024.2357569>.
- Van Veenstra, A. F., van Zoonen, E. A. & Helberger, N. (2021), ‘ELSA labs for human centric innovation in AI’, Netherlands AI Coalition. <https://nlaic.com/en/bouwsteen/human-centric-ai/elsa-concept/>.
- Vinitsky, E., Köster, R., Agapiou, J. P., Duéñez-Guzmán, E. A., Vezhnevets, A. S. & Leibo, J. Z. (2023), ‘A learning agent that acquires social norms from public sanctions in decentralized multi-agent settings’, *Collective Intelligence* **2**(2). DOI:<https://dx.doi.org/10.1177/26339137231162025>.
- Wang, C., Boerman, S. C., Kroon, A. C., Möller, J. & de Vreese, C. H. (2024), ‘The artificial intelligence divide: Who is the most vulnerable?’, *New Media & Society* **26**, 14614448241232345. DOI:<https://dx.doi.org/10.1177/14614448241232345>.

- Whitney, C. D. & Norman, J. (2024), Real risks of fake data: Synthetic data, diversity-washing and consent circumvention, *in* 'Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency', FAccT '24, Association for Computing Machinery, New York, NY, USA, p. 1733–1744. DOI:<https://dx.doi.org/10.1145/3630106.3659002>.
- Woodgate, J. & Ajmeri, N. (2022), Macro ethics for governing equitable sociotechnical systems, *in* 'Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS)', IFAAMAS, Online, pp. 1824–1828. DOI:<https://dx.doi.org/10.5555/3535850.3536118>.
- Woodgate, J. & Ajmeri, N. (2024), 'Macro ethics principles for responsible AI systems: Taxonomy and directions', *ACM Computing Surveys* **56**(11), 1–37.
- You, Z., Zhang, P., Zheng, J. & Cartledge, J. (2024), Multi-relational graph diffusion neural network with parallel retention for stock trends classification, *in* 'Proceedings of the 49th IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)', IEEE, Seoul. DOI:<https://dx.doi.org/10.1109/icassp48485.2024.10447394>.
- Young, E., Wajcman, J. & Sprejer, L. (2021), Where are the women? Mapping the gender job gap in AI, Policy briefing: Full report, The Alan Turing Institute, UK.
- Zhang, J., Wen, J. & Chen, J. (2023), 'Modelling market fluctuations under investor sentiment with a Hawkes-contact process', *The European Journal of Finance* **29**(1), 17–32. DOI:<https://dx.doi.org/10.1080/1351847X.2021.1957699>.
- Zhang, Q., Wallbridge, C. D., Jones, D. M. & Morgan, P. L. (2024), 'Public perception of autonomous vehicle capability determines judgment of blame and trust in road traffic accidents', *Transportation Research Part A: Policy and Practice* **179**, 103887. DOI:<https://dx.doi.org/10.1016/j.tra.2023.103887>.
- Zhang, T., Yang, J., Liang, N., Pitts, B. J., Prakah-Asante, K., Curry, R., Duerstock, B., Wachs, J. P. & Yu, D. (2023), 'Physiological measurements of situation awareness: A systematic review', *Human Factors* **65**(5), 737–758. DOI:<https://dx.doi.org/10.1177/0018720820969071>.
- Zhang, Y., Chapple, K., Cao, M., Dennett, A. & Smith, D. (2020), 'Visualising urban gentrification and displacement in Greater London', *Environment and Planning A: Economy and Space* **52**(5), 819–824. DOI:<https://dx.doi.org/10.1177/0308518X19880211>.