

Designing AI for Social Good: Seven Essential Factors

Josh Cowls^{1,2*}, Thomas C King¹, Mariarosaria Taddeo^{1,2}, Luciano Floridi^{1,2}

¹Digital Ethics Lab, Oxford Internet Institute, University of Oxford, Oxford, UK

²The Alan Turing Institute, London, UK

*Email of correspondence author: josh.cowls@oii.ox.ac.uk

Abstract

The idea of Artificial Intelligence for Social Good (henceforth AI4SG) is gaining traction within information societies in general and the AI community in particular. It has the potential to address social problems effectively through the development of AI-based solutions. Yet, to date, there is only limited understanding of what makes AI socially good in theory, what counts as AI4SG in practice, and how to reproduce its initial successes in terms of policies (Cath et al. 2018). This article addresses this gap by extrapolating seven ethical factors that are essential for future AI4SG initiatives from the analysis of 27 case studies of AI4SG projects. Some of these factors are almost entirely novel to AI, while the significance of other factors is heightened by the use of AI. From each of these factors, corresponding best practices are formulated which, subject to context and balance, may serve as preliminary guidelines to ensure that well-designed AI is more likely to serve the social good.

Keywords

AI4SG, Artificial Intelligence, Ethics, Social Good, Sustainable Development Goals, Privacy, Safety.

Funding

Cowls is the recipient of a Doctoral Studentship from the Alan Turing Institute. King's work was supported by a grant by Google UK Limited. Floridi's and Taddeo's work was supported by Privacy and Trust Stream - Social lead of the PETRAS Internet of Things research hub - PETRAS is funded by the Engineering and Physical Sciences Research Council (EPSRC), grant agreement no. EP/N023013/1 - and by Facebook Inc.

1. Introduction

The idea of Artificial Intelligence (AI) for Social Good (henceforth AI4SG) is becoming popular in many information societies and gaining traction within the AI community (Hager et al. 2017). Projects addressing AI4SG vary significantly. They range from models to predict septic shock (Henry et al. 2015) to game-theoretic models to prevent poaching (Fang et al. 2016); from online reinforcement learning to target HIV-education at homeless youths (Yadav, Chan, Jiang, et al. 2016) to probabilistic models to prevent harmful policing (Carton et al. 2016) and support student retention (Taddeo and Floridi 2018a). Indeed, new applications of AI4SG appear almost daily, making possible socially good outcomes that were once less easily achievable, unfeasible, or unaffordable.

Several frameworks for the design, development, and deployment of ethical AI in general have recently emerged (see Floridi et al. 2018 for a comparative analysis and synthesis). However, there is still only limited understanding about what constitutes “AI for the social good” (Taddeo and Floridi 2018a). Approaching AI4SG *ad hoc*, by analysing specific areas of application—such as famine-relief or disaster management—as an annual summit for AI industry and government has done (“AI for Good Global Summit” 2017; 2018; 2019) indicates the presence of a phenomenon, but neither explains it nor suggests how other AI4SG solutions could and should be designed to harness its full potential.

Lacking a clear understanding of what makes AI socially good in theory, what counts as AI4SG in practice, and how to reproduce its initial successes in terms of policies is a problem because designers of AI4SG face at least two main challenges: unnecessary failures and accidental successes. AI software is shaped by human values which, if not carefully selected, may lead to “good-AI-gone-bad” scenarios. For example, consider the failure of IBM’s oncology-support software, which attempts to use machine learning to identify cancerous tumours, but which was rejected by medical practitioners “on the ground” (Ross and Swetlitz 2017). The system was trained using synthetic data and was not refined enough to interpret ambiguous, nuanced, or otherwise “messy” patient health records (Strickland 2019). It also relied on US medical protocols, which are not applicable worldwide. The heedless deployment and the poor design of the software led to misdiagnoses and erroneous treatment suggestions, breaching the trust of doctors and hospitals. Context-specific design and deployment could help prevent such value misalignment and deliver successful AI4SG projects on a more consistent basis.

At the same time, the genuinely socially good outcomes of AI may arise merely by chance, for example through an accidental application of an AI solution in a different context. This was the case with the use of a different version of IBM’s cognitive system. In this case, the Watson

system was originally designed to identify biological mechanisms, but when used in a classroom setting, it inspired engineering students to solve design problems (Goel et al. 2015). In this instance, AI provided a unique mode of education. But lacking a clear understanding of AI4SG means that this success accidental; it can hardly be repeated systematically.

In order to avoid unnecessary failures and accidental successes, AI4SG would benefit from an analysis of the essential factors that underline the design of successful AI4SG systems. In this article, we provide the first, fine-grained analysis of these factors. Our aim here is not to document every single ethical consideration for an AI4SG project. For example, it is essential, and hopefully self-evident, that an AI4SG project ought not to advance the proliferation of weapons of mass destruction, an imperative which we do not discuss here (Taddeo and Floridi 2018b). Instead, we focus on factors that are *particularly relevant* to AI as a technology designed and used for the *advancement of social good*. To anticipate, these are: (1) *falsifiability and incremental deployment*; (2) *safeguards against the manipulation of predictors*; (3) *receiver-contextualised intervention*; (4) *receiver-contextualised explanation and transparent purposes*; (5) *privacy protection and data subject consent*; (6) *situational fairness*; and (7) *human-friendly semanticisation*.

The rest of the article is structured as follows. In section two, we explain how we identified the seven factors. In section three, we analyse the seven factors individually. We elucidate each of them by reference to one or more case studies, and we derive from each factor a corresponding best practice for AI4SG creators. In the concluding section, we discuss the factors and suggest how tensions between them may be resolved.

2. Methodology

AI4SG initiatives are successful insofar as they help to reduce, mitigate or eradicate a given problem of moral significance. Thus, our analysis of the essential factors for successful AI4SG is based on the following working definition:

AI4SG =_{def.} the design, development, and deployment of AI systems in ways that (i) prevent, mitigate or resolve problems adversely affecting human life and/or the wellbeing of the natural world, and/or (ii) enable socially preferable and/or environmentally sustainable (Floridi 2007) developments.

Following this definition, we analysed a set of 27 projects, obtained via desk research undertaken by the authors, by checking for clear and significant cases of successful and unsuccessful examples of AI4SG. Out of these sample, we finally identified 7 cases (see Appendix for a list), as being representative in terms of scope, variety, impact, and for their potentiality to evince shared and

constant factors that should characterise the design of AI4SG projects.¹ Of course, caution is in order when attempting to abstract lessons from case studies, especially given the diversity of contexts in which AI may be used to advance the social good. For this reason, we focused only on *essential* factors that are ethically robust, in the sense that they represent design considerations that should be ethically endorsed. Before discussing them, it is important to clarify three general features of the whole set: dependency, order, and coherence.

The seven factors are often intertwined and co-dependent, but for the sake of simplicity we discuss them separately. Nothing should be inferred from this choice. In the same way, the factors are all essential, none of them is “more important” than any other, so we shall introduce them not in terms of priority, but somewhat historically, starting with factors that pre-date AI, and yet take on greater importance when AI technologies are used, owing to the particular capabilities and risks of AI (Yang et al. 2018).² These include *falsifiability and incremental deployment* and *safeguards against the manipulation of data*. There are also factors that relate more intrinsically to the sociotechnical characteristics of AI as it exists today, like *situational fairness* and *human-friendly semanticisation*.

The factors identified in this article are coherent with more general work in the field of AI ethics. Each factor relates to at least one of five ethical principles of AI—beneficence, nonmaleficence, justice, autonomy, and explicability—identified in the comparative analysis mentioned above (Floridi et al. 2018). This coherence is crucial: AI4SG cannot be inconsistent with the ethical framework guiding the design and evaluation of AI in general. Here, the principle of beneficence is of particular relevance. It states that the use of AI should provide benefit to people and the natural world, and indeed AI4SG projects should not just respect but *reify* this principle. Beneficence is a necessary condition of AI4SG, yet it is insufficient, because the beneficent impact of an AI4SG project may be “offset” by the creation or amplification of other risks or harms.³ Ethical analysis informing the design and the deployment of AI4SG initiatives has a central role in mitigating foreseeable risks to face unintended consequences and possible misuses of the technology.

3. Seven essential factors for successful AI4SG

¹ These cases were also used as seeds for the larger survey of AI4SG projects presented in Floridi et al (2019).

² As noted in the introduction, we cannot hope to document every single ethical consideration for a social good project, so even the least novel factors here are those that take on new relevance in the context of AI.

³ This should not be taken as necessitating a utilitarian calculation: the beneficial impact of a given project may be “offset” by the violation of some categorical imperative. Therefore even if an AI4SG project would do “more good than harm”, the harm may be ethically intolerable. In such a hypothetical case, one would not be morally obliged to develop and deploy the project in question.

As we anticipated, the factors are (1) *falsifiability and incremental deployment*; (2) *safeguards against the manipulation of predictors*; (3) *receiver-contextualised intervention*; (4) *receiver-contextualised explanation and transparent purposes*; (5) *privacy protection and data subject consent*; (6) *situational fairness*; and (7) *human-friendly semanticisation*. We shall elucidate each factor with one or more examples of projects in the sample, and offer a corresponding best practice.

3.1 Falsifiability and incremental deployment

Trustworthiness is essential for technology in general (Taddeo and Floridi 2011; Taddeo 2017), and for AI4SG applications in particular, to be adopted and have a meaningful positive impact on human life and environmental wellbeing. While there is no universal rule or guideline that can ensure or guarantee trustworthiness, *falsifiability* is essential factor to improve the trustworthiness of technological applications in general, and AI4SG applications in particular.

Falsifiability entails the specification, and the possibility of empirical testing, of one or more critical requirements, that is, an essential condition, resource, or means for a capability to be fully operational, such that something could or should not work without it. *Safety* is an obvious critical requirement. Hence, for an AI4SG system to be trustworthy, its safety should be falsifiable. If falsifiability is not possible, then the critical requirements cannot be checked, and then the system should not be deemed trustworthy. This is why falsifiability is an essential factor for all conceivable AI4SG projects.

Unfortunately, we cannot know for sure that a given AI4SG application is safe unless we can test the application in all possible contexts. In this case, the map of testing would simply equate to the territory of deployment. As this *reductio ad absurdum* makes clear, complete certainty is out of reach. What is within reach, in an uncertain and fuzzy world with many unforeseen situations, is the possibility to know when a given critical requirement is not implemented or may be failing to work properly. Hence, if the critical requirements are falsifiable, we can know when the AI4SG application is not trustworthy, but not whether it is trustworthy.

Critical requirements should be tested with an incremental deployment cycle. Unintended hazardous effects may only reveal themselves after testing. At the same time, software should only be tested in the real world if it is safe to do so. This requires adoption of a deployment cycle whereby developers: (a) ensure that the application's most critical requirements or assumptions are falsifiable, (b) undertake hypothesis testing of those most critical requirements and assumptions in safe, protected contexts, and, if these hypotheses are not disproven over a small set of suitable contexts, then (c) conduct testing across increasingly wide contexts, and/or test a larger set of less-

critical requirements, and all this while (d) being ready to halt or modify the deployment as soon as hazardous or other unwanted effects may appear.

AI4SG applications may use formal approaches to try to test critical requirements. For example, they may include the use of formal verification to ensure that autonomous vehicles, and AI systems in other safety-critical contexts, would make the ethically preferable choice (Dennis et al. 2016). Such methods offer safety checks that, in terms of falsifiability, can be proved correct. Simulations may offer roughly similar guarantees. A simulation enables one to test whether critical requirements (again, consider safety) are met under a set of formal assumptions. Unlike a formal proof, a simulation cannot always indicate that the required properties are necessarily always satisfied. But a simulation often enables one to test a much wider set of cases that cannot be dealt with formally, e.g., due to the complexity of the proof.

It would be misguided to rely purely on formal properties or simulations to falsify an AI4SG application. The assumptions of these models cage the real-world applicability of any conclusions that one might make. And assumptions may be incorrect in reality. What one may prove to be correct via a formal proof, or likely correct via testing in simulation, may be disproved later with the real-world deployment of the system. For example, developers of a game-theoretic model for wildlife security assumed a relatively flat topography without serious obstructions. Hence, the software that they developed originally had an incorrect definition of an optimal patrol route. Incremental testing of the application enabled the refinement of the optimal patrol route by proving wrong the assumption of a flat topography (Fang et al. 2016).

If novel dilemmas in real-world contexts require the alteration of prior assumptions made in the lab, one solution is to rectify *a priori* assumptions after deployment. Alternatively, one may adopt an “on-the-fly” or runtime system for a constant update of a program’s processing (“understanding”) of its inputs. Yet, problems also abound with this approach. For example, Microsoft’s infamous Twitter bot, Tay, acquired meanings, in a very loose sense, at runtime, as it learned from Twitter users how it should respond to tweets. After deployment in the real—and frequently vicious—world of social media, however, the bot’s ability to adapt constantly its “conceptual understanding” became an unfortunate bug, as Tay “learned” and regurgitated offensive language and unethical associations between concepts from other users (Neff and Nagy 2016).

The use of a retrodictive approach—that is, an attempt to understand some aspect of reality through *a priori* information—to deal with the falsifiability of requirements presents similar problems. This is noteworthy, since retrodiction is the primary method of supervised machine

learning approaches that learn from data (e.g., the learning of a continuous transformation function in the case of neural networks).

From the previous analysis it follows that the essential factor of falsifiability and incremental development comprises a cycle: engineering requirements that are falsifiable (so that it is at least possible to know whether the requirements are not met); falsification testing for incrementally improving levels of trustworthiness; adjustment of *a priori* assumptions; and then and only then deployment in an incrementally wider and critical context. Germany's approach to regulating autonomous vehicles offer a good example of this incremental approach. Deregulated zones allow experimentation of constrained autonomy and, after increasing the levels of trustworthiness, manufacturers may test vehicles with higher levels of autonomy (Pagallo 2017). Indeed, the creation of such deregulated zones, or *teststrecken*, was one recommendation to support more ethical AI policy at the European level (Floridi et al. 2018). The identification of this essential factor yields the following best practice:

- 1) *AI4SG designers should identify falsifiable requirements and test them in incremental steps from the lab to the "outside world".*

3.2 Safeguards against the manipulation of predictors

The use of AI to predict future trends or patterns is very popular in AI4SG contexts, from applying automated prediction to redress academic failure (Lakkaraju et al. 2015), to preventing illegal policing (Carton et al. 2016), and detecting corporate fraud (Zhou and Kapoor 2011). The predictive power of AI4SG faces two risks: the manipulation of input data, and excessive reliance on non-causal indicators.

The manipulation of data is not a new problem, nor is it limited to AI systems alone. But AI may exacerbate it, and it is a noteworthy risk for any AI4SG initiative, because it can impair the predictive power of AI and lead to the avoidance of socially good interventions at the individual level. Consider the concern raised by Ghani over teachers who face being evaluated in respect to: the percentage of students in their class who are above a certain risk threshold. If the model was transparent—for example, heavily reliant on math GPA—the teacher could inflate math grades and reduce the intermediate risk scores of their students (Ghani 2016).

As Ghani goes on to argue, the same concern applies to predictors of adverse police officer interactions:

these systems [are] very easy to understand and interpret, but that also makes them easy to game. An officer who has had two uses of force in the past 80 days may choose to be a bit more careful over the next 10 days, until the count rolls over to zero again.

These hypothetical examples make clear that, when the model used is an easy one to understand “on the ground”, it is already open to abuse or “gaming”, independently of whether AI is used. The introduction of AI complicates matters, owing to the scale at which AI is typically applied. As we have seen, if the information used to predict a given outcome is known, an agent with such information (that is predicted to take a particular action) can change each predictive variable’s value in order to avoid an intervention. In this way, the predictive power of the overall model is reduced, as has been shown by empirical research in the domain of corporate fraud (Zhou and Kapoor 2011). Such a phenomenon could carry across from fraud detection to the domains that AI4SG initiatives seek to address.⁴

At the same time, there is a risk that excessive reliance on non-causal indicators – that is, data which is correlated with, but not causal of, a phenomenon – may distract attention from the context in which the AI4SG designer is seeking to intervene. To be effective, any such intervention should alter the underlying causes of a given problem, such as a student’s domestic problems or inadequate corporate governance, rather than non-causal predictors. To do otherwise is to risk addressing only a symptom, rather than the root cause of a problem.

These risks suggest the need to consider the use of safeguards as a design factor for AI4SG projects. Such safeguards may constrain the selection of indicators to be used in the design of AI4SG projects; the extent to which these indicators should shape interventions; and/or the level of transparency that should apply to how indicators affect decision. This yields the following best practice:

- 2) *AI4SG designers should adopt safeguards which (i) ensure that non-causal indicators do not inappropriately skew interventions, and (ii) limit, when appropriate, knowledge of how inputs affect outputs from AI4SG systems, to prevent manipulation.*

3.3 Receiver-contextualised intervention

It is essential that software intervenes in users’ life only in ways that respects their autonomy. Again, this is not a problem that arises only with AI-driven interventions, but the use of AI introduces new considerations. In particular, a core challenge for AI4SG projects is to devise interventions that balance current and future benefits. The balancing problem, which is familiar to preference-elicitation research (Boutilier 2002; Faltings et al. 2004; Chajewska, Koller, and Parr 2000), boils down to a temporal choice interdependency. An intervention in the present can elicit user preferences that then enable the software to contextualise future interventions to the given user. Consequently, an intervention strategy that has no impact on user autonomy (e.g., one that

⁴ For a discussion of the use of artificial intelligence in criminal acts more generally, see King et al. 2019).

lacks any interventions) may be ineffective in extracting the necessary information for correctly contextualised future interventions. Conversely, an intervention that overly infringes upon a user's autonomy may cause the user to reject the technology, making future interventions impossible.

This balancing consideration is a common one for AI4SG initiatives. Take, for example, interactive activity recognition software for people with cognitive disabilities (Chu et al. 2012). The software is designed to prompt patients to maintain a daily schedule of activities (e.g., taking medication), whilst minimising interruptions to their wider goals. Each intervention is contextualised in such a way that the software learns the timing of future interventions from responses to past interventions. Moreover, only important interventions are made, and yet all interventions are partially optional because declining one prompt leads to the same prompt later on. Here, the concern was that patients would reject an overly intrusive technology; hence a balance was sought. This balance is lacking in our second example. A game-theoretic application intervenes in wildlife security officers' patrols by offering suggested routes (Fang et al. 2016). If a route poses physical obstacles, however, then the software lacks the possibility to provide alternative suggestions. Officers may ignore the advice by taking a different route, but not without disengaging from the application. It is essential to relax such constraints, so that users can ignore an intervention, but accept subsequent, more appropriate interventions (in the form of advice) later on.

These examples point to the importance of seeing users as equal partners in both the design and deployment of autonomous decision-making systems. The adoption of this mindset might have helped prevent the tragic loss of two Boeing 737 Max airliners. It appears that the pilots of these flights struggled to reverse a software malfunction caused by faulty sensors, due in part to the absence of "optional safety features" which Boeing sold separately (Tabuchi and Gelles 2019).

The risk of false positives (unnecessary intervention, creating disillusionment) is often just as problematic as false negatives (no intervention where it is necessary, limiting effectiveness). Hence, a suitable receiver-contextualised intervention is one that achieves the right level of disruption while respecting autonomy through optionality. This contextualisation rests on information about users' capacities, preferences and goals, and the circumstances in which the intervention will take effect.

One can consider five dimensions relevant to a receiver-contextualised intervention. Four of these dimensions emerge from McFarlane's taxonomy of interdisciplinary research on disruptive computer-human interruptions (McFarlane 1999; McFarlane and Latorella 2002 17-19). These are: the individual characteristics of the person receiving the intervention; the methods of coordination between the receiver and the system; the meaning or purpose of the intervention; and the overall

effects of the intervention.⁵ A fifth dimension of relevance is optionality – a user can choose either to ignore all offered advice or to drive the process and request a different intervention better suited to their needs.

We can summarise these five dimensions in the form of the following best practice for receiver-contextualised intervention:

- 3) *AI4SG designers should build decision-making systems in consultation with users interacting with, and impacted, by these systems; with understanding of users' characteristics, of the methods of coordination, and the purposes and effects of an intervention; and with respect for users' right to ignore or modify interventions.*

3.4 Receiver-contextualised explanation and transparent purposes

AI4SG applications should be designed to make explainable the operations and outcomes of these systems and to make transparent their purposes. These two requirements are of course intrinsically linked, as the operations and outcomes of AI systems reflect the wider purposes of human designers; in this section, we address both in turn.

Making AI systems explainable is an important ethical principle (Floridi et al. 2018). It has been a focus of research since at least 1975 (Shortliffe and Buchanan 1975). And it has gained more attention recently (Thelisson, Padh, and Celis 2017; Wachter, Mittelstadt, and Floridi 2016) given the increasingly pervasive distribution of AI systems. As we saw above, AI4SG projects should offer interventions that are contextualised to the receiver. In addition, the *explanation* for an intervention should be contextualised in order to be adequate.

Designers of AI4SG projects have tried to increase the explainability of decision-making systems in various ways. For example, researchers have used machine learning to predict academic adversity (Lakkaraju et al. 2015). These predictors used concepts that the school officials interpreting the system found familiar and salient, such as GPA scores and socio-economic categorisations. Researchers have also used reinforcement-learning to help officials at homeless shelters educate homeless youths about HIV (Yadav, Chan, Xin Jiang, et al. 2016). The system learns how to maximise the influence of HIV education, by choosing which homeless youths to educate, on the basis that homeless youths may pass on their knowledge. One version of the system explained which youth was chosen by revealing their social network graph. However, the homeless shelter officials found that these explanations were counter-intuitive, potentially affecting the

⁵The four remaining dimensions proposed by Macfarlane — the source of the interruption, the method of expression, the channel of conveyance and the human activity changed by the interruption — are not relevant for purpose of this article.

understanding of how the system worked and, hence, users' trust in the system. These two cases exemplify the importance of the right conceptualisation when explaining an AI-based decision.

The right conceptualisation is likely to vary between AI4SG projects, because they differ greatly in their objectives, subject matter, context and stakeholders. The conceptual framework, that is, the Level of Abstraction (LoA; Floridi 2017), depends on what is being explained and to whom. A LoA is a key component of a theory, and hence of any explanation. A theory comprises five components:

1. a *System*, which is the referent or object analysed by a theory;
2. a *Purpose*, which is the “what for” that motivates the analysis of a system (note that this answers the question “what is the analysis for?” and should not be confused with a system’s purpose, which answers the question “what is the system for?”. Below, we use the term “goal” for system’s purpose whenever there may be a risk of confusion);
3. a *Level of Abstraction*, which provides a lens through which a system is analysed, and generates;
4. a *Model*, that is, some relevant a reliable information about the analysed system, which identifies;
5. a *Structure* of the system, which comprises the features that belong to the system being analysed.

There is an interdependency between the choice of the specific purpose, the relevant LoA that can fulfil the purpose, the system analysed, and the model obtained by analysing the system at a specified LoA for a particular purpose. The LoA provides the conceptualisation of the system (e.g., GPA scores, and socio-economic backgrounds). But the purpose constrains the construction of LoAs. For example, if we choose to explain the decision making system itself (e.g., the use of particular machine learning techniques), then the LoA can only conceptualise those AI techniques. In turn, the LoA generates the model, which explains the system. The model identifies system structures, such as a specific student’s GPA score, poor attendance rate, and their socioeconomic background being predictors of their academic failure. Consequently, designers must choose carefully the purpose and the corresponding LoA, so that the explanation model can provide the right explanation of the system in question for a given receiver.

A LoA is chosen for a specific *purpose*: for example, a LoA chosen to explain a decision taken on the basis of outcomes obtained through an algorithmic procedure varies depending on whether the explanation is meant for the receiver of that decision or for an engineer responsible for the design of the algorithmic procedure. This is because, depending on the purpose and its granularity (e.g. a customer-friendly vs. engineer-friendly explanation), not every LoA is

appropriate for a given receiver. Sometimes, a receiver's conceptual view of the world may differ from the one on which the explanation is based. In other cases, a receiver and an explanation may be conceptually aligned, but the receiver may not agree on the level of granularity (LoA) of the information (what we called more precisely the model) provided. Conceptual disalignment means that the receiver may find the explanation irrelevant, unintelligible or, as we shall see below, questionable. In respect of (un)intelligibility, a LoA may use unknown labels (so-called observables), or labels that have different meanings for different users.

Empirical studies (Gregor and Benbasat 1999) suggest that the suitability of an explanation differs among receivers according to their expertise. Receivers may require explanations about how the AI software came to a decision, especially when they must take action based on that decision (Gregor and Benbasat 1999; Watson et al. 2019). How the AI system came to a conclusion can be just as important as the justification for that conclusion. Consequently, designers must also contextualise the method of explanation to the receiver.

The case of the software that uses influence-maximisation algorithms to target homeless youths for HIV education provides a good example of the relevance of the receiver-contextualisation of concepts (Yadav, Chan, Jiang, et al. 2016). The researchers involved in this project considered three possible LoAs when designing the explanation model: the first LoA included utility calculations; the second LoA focused on social graph connectivity; and a third LoA focusing on pedagogic purpose. The first LoA highlighted the utility of targeting one homeless youth over another. According to the researchers, in this case, homeless shelter workers (the receivers) might have misunderstood the utility calculations or found them irrelevant. Utility calculations offer little explanatory power beyond the decision itself, because they often simply show that the "best" choice was made, and how good it was. Explanations based on the second LoA faced a different problem: the receivers assumed that the most central nodes in the network were the best for maximising the influence of education, while the optimal choice is often a set of less well-connected nodes. The third LoA was eventually chosen, after subsequent user testing of different explanation frameworks (Yadav, Chan, Xin Jiang, et al. 2016). This LoA led the researchers to introduce a comparison between the optimal strategy that the system would choose, and the one of the officials. This provided users with the most important information on why their choice would be suboptimal.

Given a particular system, what purpose one chooses to pursue when seeking an explanation of it, at what LoA, and what issuing model is obtained are crucial variables that impact the effectiveness of an explanation. Explainability breeds trust in, and fosters adoption of, AI4SG solutions (Herlocker, Konstan, and Riedl 2000; Swearingen and Sinha 2002; Bilgic and Mooney

2005). This is why it is essential that software uses persuasive *argumentation* for the target audience. This is likely to include information about both the general functionality and logic employed by a system and the reasons for the specific decision being made (Wachter, Mittelstadt, and Floridi 2017).

Transparency in the goal (i.e., system's purpose) of the system is also crucial. Consider, for example, the development of AI solutions to prompt people with cognitive disabilities to take their medication (Chu et al. 2012). On its face, this application may seem invasive, involving a vulnerable user, limiting the effectiveness of receiver-conceptualised explanation. However, the system is not designed to coerce the patients into a given behaviour, nor is it designed to resemble a human being. The patients have autonomy not to interact with the AI system in question. This case highlights the importance of transparency in goals, particularly in contexts in which explainable operations and outcomes are unworkable or undesirable. Transparency in goals, thus, undergirds other safeguards around the protection of target populations and may help ensure compliance with relevant legislation and precedent (Reed 2018).

Conversely, opaque goals may prompt misunderstanding and the potential for harms. For instance, when users of an AI system are unclear about what type of agent they are dealing with—human, artificial, or a hybrid combination of both—they may wrongly assume that the tacit norms of human-to-human social interaction are upheld (e.g., not recording every detail of a conversation) (R. Kerr 2003). As ever, the social context in which an AI4SG application takes place impacts the extent to which AI systems should be transparent in their operations. Because transparency is the default but not absolute position, there may be valid reasons for designers to obviate informing users of the software's goals. For example, the scientific value of a project or the health and safety conditions of a public space may justify *temporarily* opaque goals. Consider a study that deceived students into believing that they were interacting with a human course-assistant that was in fact, over time, realised to be a bot (Eicher, Polepeddi, and Goel 2017). The bot's deception, as the authors argue, was for playing the "imitation game" without causing the students to choose simpler and less human-like natural-language queries based on preconceptions of AI capabilities. In such cases, the choice between opacity and transparency may be informed by preexisting notions of informed consent for human-subject experiments embedded in the Nuremberg Code, The Declaration of Helsinki, and The Belmont Report (Nijhawan et al. 2013).

More broadly, the ability to avoid the use of an AI system becomes more likely when AI software reveals its endogenous goals, like classifying data about a person. For example, AI software could inform staff in a hospital ward that it has the goal of classifying their hygiene levels (Haque et al. 2017). In this case, the staff may decide to avoid such classifications if there are

reasonable alternative actions that they can take. In other cases, revealing a goal makes it less likely to be fulfilled.

Making transparent the goals and motivations of AI4SG developers themselves is an essential factor to the success of any project, but one that may contrast the very purpose of the system. This is why it is crucial to assess, at a design stage, what level of transparency (i.e. how much transparency, of what kind, for whom, and about what?) the project will embrace given its overall goal and the context of implementation. Taken together with the need for receiver-conceptualised explanation, this consideration yields the following set of best practices:

- 4) *AI4SG designers should choose a Level of Abstraction for AI explanation that fulfils the desired explanatory purpose and is appropriate to the system and the receivers; then deploy arguments that are rationally and suitably persuasive for the receivers to deliver the explanation; and ensure that the goal (the system's purpose) for which an AI4SG system is developed and deployed is knowable to receivers of its outputs by default.*

3.5 Privacy protection and data subject consent

Of our seven factors, privacy has perhaps the most voluminous literature. This should not be a surprise, since privacy is considered to be an essential condition for safety, human dignity, and social cohesion, among other things (Solove 2008), and because earlier waves of digital technology have already had a major impact on privacy (Nissenbaum 2009). People's safety may be compromised when a state gains control over individuals via privacy infringements (Taddeo 2014; Lynskey 2015). Respect for privacy is also a necessary condition of human dignity: since we can view personal information as constituting an individual, and deprivatising records without consent is likely to constitute a violation of human dignity (Floridi 2016). The conception of individual privacy as a fundamental right underlies legislative action in Europe and judicial decisions in India (Mohanty and Bhatia 2017). Privacy supports people in deviating from social norms without causing offense, and communities in maintaining their social structures, so privacy undergrids also social cohesion.

In the case of AI4SG, it is particularly important to emphasise the relevance of users' consent to the use of personal data. Tensions may arise between different thresholds of consent (Price and Cohen 2019). The tension is often at its most fraught in "life-or-death" situations such as national emergencies and pandemics. Consider the outbreak of Ebola in West Africa in 2014, which posed a complex ethical dilemma (*The Economist* 2014). In this case, the rapid release and analysis of call-data records from cell phone users in the region may have allowed epidemiologists

to track the spread of the deadly disease. However, the release of the data was held up over valid concerns around users' privacy, as well as the value of the data to industrial competitors.

In circumstances where haste is not so crucial, it is possible to obtain a subject's consent for – and before – the data being used. The level or type of consent sought can vary with the context. In healthcare, one may adopt an assumed consent threshold, whereby reporting a medical issue to a doctor constitutes assumed consent on the part of a patient. In other circumstances, an informed consent threshold will be more appropriate. Yet, since informed consent requires researchers to obtain a patient's specific consent before using their data for a non-consented purpose, practitioners may choose an explicit consent threshold to general data processing, i.e., for any medical usage. This threshold does not require informing the patient about all of the possible ways that researchers may use their data (Etzioni 1999). Another alternative is the evolving notion of “dynamic consent”, whereby individuals can monitor and adjust their privacy preferences on a granular level (Kaye et al. 2015)

In other cases, informed consent may be waived altogether. This was the case with the recent creation of machine learning software to predict the prognosis of ovarian cancer sufferers by drawing upon retrospective analysis of anonymised images (Lu et al. 2019). The use of patient health data in the development of AI solutions without patients' consent has also attracted the attention of data protection regulators. In 2017, the UK's Information Commissioner ruled that the Royal Free NHS Foundation Trust violated the Data Protection Act when it provided patient details to Google DeepMind, for the purposes of training an AI system to diagnose acute kidney injury (Burgess 2017). The Commissioner noted as a “shortcoming” that “patients were not adequately informed that their data would be used as part of the test” (“Royal Free - Google DeepMind Trial Failed to Comply with Data Protection Law” 2017).

Striking a balance between respecting patient privacy and creating effective AI4SG is still possible, however. This was the challenge faced by the researchers in Haque et al. (2017), who wanted to create a system for tracking compliance with rules around hand hygiene in hospitals, to prevent the spread of infections. Despite the clear technical advantages of taking a computer vision-based approach to the problem, the use of video recording runs up against privacy regulations constraining it. Even in cases where video recording is allowed, access to the recordings (in order to train an algorithm) is often strict. Instead, the researchers resorted to “depth images”, which de-identify subjects, preserving their privacy. While this design choice meant “losing important visual appearance cues in the process” (3), it satisfied privacy rules, and the researchers' non-intrusive system still managed to outperform existing solutions.

Finally, consent in the online space is also problematic, users often lack the choice and are presented with a ‘take it or leave it’ option when accessing online services (Nissenbaum 2011; Taddeo and Floridi 2015). The relative lack of protection or consent for the second-hand use of personal data that is publicly shared online enables the development of ethically problematic AI software. For example, a recent paper used publicly available images of faces uploaded to a dating website as a way to train AI software to detect someone’s sexuality based on a small number of photos (Wang and Kosinski 2018). While the study received ethics committee approval, it raises further questions around consent, since it is implausible that the users of the dating website could or necessarily would have consented to the use of their data for this particular purpose.

Privacy is not a novel problem, but the centrality of personal data to many AI (and AI4SG) applications heightens its ethical significance and creates issues around consent (Taddeo and Floridi 2018a). From this we can derive the following best practice:

- 5) *AI4SG designers should respect the threshold of consent established for the processing of datasets of personal data.*

3.6 Situational fairness

AI developers typically rely on data, which may be biased in ways that are socially significant. This bias may carry across to the algorithmic decision-making that underpins many AI systems, in ways that are unfair to the subjects of the decision-making process (Caliskan, Bryson, and Narayanan 2017). These decision may be based on factors of ethical importance (e.g., ethnic, gender, or religious grounds) and irrelevant to the decision-making at hand, or they may be relevant but legally protected as a nondiscriminatory characteristic (Friedman and Nissenbaum 1996). Moreover, AI-driven decisions may be amalgamated from factors that are not of obvious ethical importance, and yet collectively constitute unfairly biased decision-making (Pedreshi, Ruggieri, and Turini 2008; Floridi 2012).

AI4SG initiatives relying on biased data may propagate this bias through a vicious cycle (Yang et al. 2018). Such a cycle would begin with a biased dataset informing a first phase of AI decision-making, resulting in discriminatory actions, leading to the collection and use of biased data in turn. Consider the use of AI to predict preterm birth in the United States, where the health outcomes of pregnant women have long been affected by their ethnicity. Longstanding bias against African-American women seeking treatment, owing to harmful historical stereotypes, contributes to a maternal morbidity rate that is over three times higher than that of white women (CDC 2019). Here, AI may offer great potential to reduce this stark racial divide, but only if the same historical discrimination is not replicated in AI systems (Banjo 2018). Or consider the use of predictive

policing software. Developers may train predictive policing software on policing data that contains deeply ingrained prejudices. When discrimination affects arrest rates, it becomes embedded in prosecution data (Lum and Isaac 2016). Such biases may cause discriminatory decisions (e.g., warnings or arrests) that feed back into the increasingly biased datasets (Crawford 2016), thereby completing a vicious cycle.

These examples involve the use of AI to improve outcomes in domains where data were already collected. Yet, in many other contexts, AI4SG projects (or indeed similar initiatives) are, in effect, making citizens “visible” in ways that they previously were not, including in global South contexts (Taylor and Broeders 2015). This increased visibility stresses the importance of protecting against the potential amplification of harmful bias by AI technologies.

Clearly, designers must sanitise the datasets used to train AI. However, there is equally a risk of applying too strong a disinfectant, so to speak, by removing important contextual nuances which could improve ethical decision-making. So, designers must also ensure that AI decision-making maintains sensitivity to factors that are important for inclusiveness. For instance, we should ensure that a word processor interacts identically with a human user regardless of that user’s gender and ethnicity, but also expect that it may operate in a non-equal and yet equitable way by aiding people with visual impairments.

Such expectations are not always met in the context of AI-driven reasoning. Compared to the word processor, AI makes possible a far wider range of decision-making and interaction modalities, many of which are driven by potentially biased data. Training datasets may contain natural language that carries unfair associations between genders and words which, in turn, carry normative power (Caliskan, Bryson, and Narayanan 2017). In other contexts and use cases, an equitable approach may *require* differences in communication, based on factors such as gender. Consider the case of the virtual teaching assistant which *failed* to discriminate sufficiently well between men and women in its responses to being told that a user was expecting a baby, congratulating the men and ignoring the women (Eicher, Polepeddi, and Goel 2017). A BBC News investigation highlighted an even more egregious example: a mental health chatbot deemed suitable for use by children was unable to understand a child explicitly reporting underage sexual abuse (White 2018). As these cases make clear, the use of AI in human-computer interactions, such as chatbots, requires the correct understanding of both the salient groups to which a user belongs and the characteristics they embody when they interact with the software.

Respecting situational fairness is essential for the successful implementation of AI4SG. To achieve it, AI4SG projects need to remove factors (and their proxies) that are of ethical importance but irrelevant to an outcome, and include the same factors when these are required, whether for

the sake of inclusiveness, safety, or other ethical considerations. The problem of historical biases affecting future decision-making is an old one. What is new is the potential that these biases will be embedded in, strengthened, and perpetuated anew by erroneous reinforcement learning mechanisms. This risk is especially pronounced when considered alongside the risk of opacity in AI decision-making systems and their outcomes. We will return to this topic in the next section.

From our identification of situational fairness as an essential factor, we can yield the following best practice:

- 6) *AI4SG designers should remove from relevant datasets variables and proxies that are irrelevant to an outcome, except when their inclusion supports inclusivity, safety, or other ethical imperatives.*

3.7 Human-friendly semanticisation

AI4SG must allow humans to curate and foster their “semantic capital”, that is,

any content that can enhance someone’s power to give meaning to and make sense of (*semanticise*) something (Floridi 2018).

With AI, we may often have the technical capacity to automate meaning- and sense-creation (semanticisation), but mistrust or unfairness may also arise if we do so carelessly. Two problems emerge. The first problem is that AI software may define semanticisation in a way that diverges from our own choices. This is the case if a procedure arbitrarily defines meanings (e.g., based on a coin toss). The same problem may arise if AI software support some kind of semanticisation based on preexisting uses. For example, researchers have developed an application that *predicts the legal meaning* of ‘violation’ based on past cases (Al-Abdulkarim, Atkinson, and Bench-Capon 2015). If one used the software to *define* the meaning of ‘violation’,⁶ then one would end up limiting the role of judges and justices. They would no longer be able to semanticise (refine and re-define the meaning, and the possibility of making sense of) “violation”, when they interpret the law. This is a problem, because past usage does not always predict how we would semanticise the same concepts or phenomena in the future.

The second problem is that, in a social setting, it would be impractical for AI software to define all meanings and senses. Some semanticisation is subjective, because who or what is involved in the semanticisation is also partly constitutive of the process and its outcome. For example, only legally empowered agents can define the legal meaning of ‘violation’. Likewise, the meaning and sense of affective symbols, such as facial expressions, also depends on the type of agent showing a given expression. Affective AI can detect an emotion (Martínez-Miranda and

⁶ There is no suggestion that this is the intended use.

Aldea 2005), an artificial agent may state accurately that a human *appears sad*, but cannot change the meaning of sadness.

The solution to these two problems rest on distinguishing between tasks that should and should not be delegated to an artificial system. AI should be deployed to *facilitate* human-friendly semanticisation, but not to provide it itself. This is true, for example, when considering patients with Alzheimer’s disease. Research into carer-patient relations highlights three points (Burns and Rabins 2000). First, carers play a critical, but burdensome, role in reminding patients of the activities in which they participate, e.g., taking medication. Second, carers also play a critical role in providing patients with meaningful interaction. And third, when carers remind patients to take their medication, the patient-carer relation may become weaker by annoying the patient, with the carer losing some capacity to provide empathy and meaningful support. Consequently, researchers have developed AI software that balances reminding the patient against annoying the patient (Chu et al. 2012). The balance is learned and optimised using reinforcement learning. The researchers designed the system so that caregivers can spend most of their time providing empathic support and preserving a meaningful relationship with the patient. As this example shows, it is possible to use AI to sweep away formulaic tasks whilst sustaining human-friendly semanticisation.

Human-centric semanticisation, as an essential factor for AI4SG, underpins our final best practice:

- 7) *AI4SG designers should not hinder the ability for people to semanticise (that is, to give meaning to, and make sense of) something.*

4. Conclusion: Balancing Factors for AI for Social Good

The seven factors analysed in the previous pages are summarised in the following Table 1, together with the corresponding best practices.

Factors	Corresponding best practices
Falsifiability and incremental deployment	Identify falsifiable requirements and test them in incremental steps from the lab to the “outside world”.
Safeguards against the manipulation of predictors	Adopt safeguards which (i) ensure that non-causal indicators do not inappropriately skew interventions, and (ii) limit, when appropriate, knowledge of how inputs affect outputs from AI4SG systems, to prevent manipulation.
Receiver-contextualised intervention	Build decision-making systems in consultation with users interacting with and impacted by these systems; with understanding of users’ characteristics, the methods of coordination, the purposes and effects of an intervention; and with respect for users’ right to ignore or modify interventions.

Receiver-contextualised explanation and transparent purposes	Choose a Level of Abstraction for AI explanation that fulfils the desired explanatory purpose and is appropriate to the system and the receivers; then deploy arguments that are rationally and suitably persuasive for the receiver to deliver the explanation; and ensure that the goal (the system’s purpose) for which an AI4SG system is developed and deployed is knowable to receivers of its outputs by default.
Privacy protection and data subject consent	Respect the threshold of consent established for the processing of datasets of personal data.
Situational fairness	Remove from relevant datasets variables and proxies that are irrelevant to an outcome, except when their inclusion supports inclusivity, safety, or other ethical imperatives.
Human-friendly semanticisation	Do not hinder the ability for people to semanticise (that is, to give meaning to, and make sense of) something.

Table 1: Summary of seven factors supporting AI4SG and the corresponding best practices.

The seven factors suggest that creating successful AI4SG requires two kinds of balances to be struck: *intra* and *inter*.

On the one hand, each single factor in and of itself may require an intrinsic balance, for example, between the risk of over-intervening and the risk of under-intervening when devising contextual interventions; or between protection-by-obfuscation and protection-by-enumeration of salient differences between people, depending on the purposes and context of a system. On the other hand, balances are not just specific to a single factor; they are also systemic, because they must also be struck between multiple factors. Consider the tension between preventing malicious actors from understanding how to “game” the input data of AI prediction systems versus enabling humans to override genuinely flawed outcomes; or the tension between ensuring the effective disclosure of the reasons behind a decision without compromising the consensual anonymity of data subjects.

The overarching question facing the AI4SG community is, for each given case, whether one is morally obliged to, or obliged not to, design, develop, and deploy a specific AI4SG project. This article does not seek to answer such a question in the abstract. Resolving the tensions that are likely to arise among and between factors is highly context-dependent, and the previous analysis is not meant to cover all potential contexts, not least because this would be inconsistent with the argument for falsifiable hypothesis testing and incremental deployment supported in this article; nor would a checklist of purely technical “dos and don’ts” suffice. Rather, this article offers a framework of essential factors that need to be considered, interpreted and evaluated contextually when one is designing, developing, and deploying a specific AI4SG project. The future of AI4SG

will likely provide more opportunities to enrich such a framework of essential factors. AI itself may help to manage its own life cycle by providing, in a meta-reflective way, tools to evaluate how best to strike the individual and systemic balances indicated above. This article seeks to lay the ground both for good practices and policies and for further research on the ethical considerations that should undergird AI4SG projects, and hence *the* “AI4SG project” at large.

The subsequent questions of *how* these factors should be evaluated and resolved, *by whom*, and with *what supporting mechanism* (e.g. regulation or voluntary codes of conduct) are intertwined with wider ethical and political challenges regarding who has the power to do so. In other words, this concerns what legitimates decision-making with and about AI, and it is one that we leave to future research.

References

- “AI for Good Global Summit - 28-31 May 2019, Geneva, Switzerland.” n.d. AI for Good Global Summit. Accessed April 12, 2019. <https://aiforgood.itu.int/>.
- Al-Abdulkarim, Latifa, Katie Atkinson, and Trevor Bench-Capon. 2015. “Factors, Issues and Values: Revisiting Reasoning with Cases.” In *Proceedings of the 15th International Conference on Artificial Intelligence and Law*, 3–12. ICAIL ’15. New York, NY, USA: ACM. <https://doi.org/10.1145/2746090.2746103>.
- Banjo, Omotayo. 2018. “Bias In Maternal AI Could Hurt Expectant Black Mothers.” *Medium* (blog). September 21, 2018. <https://medium.com/theplug/bias-in-maternal-ai-could-hurt-expectant-black-mothers-e41893438da6>.
- Bilgic, Mustafa, and Raymond Mooney. 2005. “Explaining Recommendations: Satisfaction vs. Promotion.” In .
- Boutillier, Craig. 2002. “A POMDP Formulation of Preference Elicitation Problems.” *Proceedings of the National Conference on Artificial Intelligence*, May.
- Burgess, Matt. 2017. “NHS DeepMind Deal Broke Data Protection Law, Regulator Rules.” *Wired UK*, July 3, 2017. <https://www.wired.co.uk/article/google-deepmind-nhs-royal-free-ico-ruling>.
- Burns, Alistair, and Peter Rabins. 2000. “Carer Burden in Dementia.” *International Journal of Geriatric Psychiatry* 15 (S1): S9–S13.
- Caliskan, Aylin, Joanna J. Bryson, and Arvind Narayanan. 2017. “Semantics Derived Automatically from Language Corpora Contain Human-like Biases.” *Science* 356 (6334): 183–86. <https://doi.org/10.1126/science.aal4230>.
- Carton, Samuel, Jennifer Helsby, Kenneth Joseph, Ayesha Mahmud, Youngsoo Park, Joe Walsh, Crystal Cody, CPT Estella Patterson, Lauren Haynes, and Rayid Ghani. 2016. “Identifying Police Officers at Risk of Adverse Events.” In *Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 67–76. KDD ’16. New York, NY, USA: ACM. <https://doi.org/10.1145/2939672.2939698>.
- Cath, Corinne, Sandra Wachter, Brent Mittelstadt, Mariarosaria Taddeo, and Luciano Floridi. 2018. “Artificial Intelligence and the ‘Good Society’: The US, EU, and UK Approach.” *Science and Engineering Ethics* 24 (2): 505–528.
- Chajewska, Urszula, Daphne Koller, and Ronald Parr. 2000. “Making Rational Decisions Using Adaptive Utility Elicitation.” In *AAAI/IAAI*, 363–369.

- Chu, Yi, Young Chol Song, Richard Levinson, and Henry Kautz. 2012. “Interactive Activity Recognition and Prompting to Assist People with Cognitive Disabilities.” *Journal of Ambient Intelligence and Smart Environments* 4 (5): 443–59. <https://doi.org/10.3233/AIS-2012-0168>.
- Crawford, Kate. 2016. “Artificial Intelligence’s White Guy Problem.” *The New York Times*. June 25, 2016. <https://www.nytimes.com/2016/06/26/opinion/sunday/artificial-intelligences-white-guy-problem.html>.
- Dennis, Louise, Michael Fisher, Marija Slavkovic, and Matt Webster. 2016. “Formal Verification of Ethical Choices in Autonomous Systems.” *Robotics and Autonomous Systems* 77 (March): 1–14. <https://doi.org/10.1016/j.robot.2015.11.012>.
- Eicher, Bobbie, Lalith Polepeddi, and Ashok Goel. 2017. “Jill Watson Doesn’t Care If You’re Pregnant: Grounding AI Ethics in Empirical Studies.” In *AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society, New Orleans, LA*. Vol. 7.
- Etzioni, Amitai. 1999. “Enhancing Privacy, Preserving the Common Good.” *Hastings Center Report* 29 (2): 14–23.
- Faltings, Boi, Pearl Pu, Marc Torrens, and Paolo Viappiani. 2004. “Designing Example-Critiquing Interaction.” In *Proceedings of the 9th International Conference on Intelligent User Interfaces, 22–29. IUI ’04*. New York, NY, USA: ACM. <https://doi.org/10.1145/964442.964449>.
- Fang, Fei, Thanh H. Nguyen, Rob Pickles, Wai Y. Lam, Gopalasamy R. Clements, Bo An, Amandeep Singh, Milind Tambe, and Andrew Lemieux. 2016. “Deploying PAWS: Field Optimization of the Protection Assistant for Wildlife Security.” In *Twenty-Eighth IAAI Conference*. <https://www.aaai.org/ocs/index.php/IAAI/IAAI16/paper/view/11814>.
- Floridi, Luciano. 2007. “Global Information Ethics.” *International Journal of Technology* 3 (3): 1–11. <https://doi.org/DOI:10.4018/jthi.2007070101>.
- . 2012. “Distributed Morality in an Information Society.” *Science and Engineering Ethics* 19 (3): 727–43. <https://doi.org/10.1007/s11948-012-9413-4>.
- . 2016. “On Human Dignity as a Foundation for the Right to Privacy.” *Philosophy & Technology* 29 (4): 307–12. <https://doi.org/10.1007/s13347-016-0220-8>.
- . 2017. “The Logic of Design as a Conceptual Logic of Information.” *Minds and Machines*. <https://doi.org/10.1007/s11023-017-9438-1>.
- Floridi, Luciano, Josh Cowsls, Monica Beltrametti, Raja Chatila, Patrice Chazerand, Virginia Dignum, Christoph Luetge, Robert Madelin, Ugo Pagallo, and Francesca Rossi. 2018. “AI4People—An Ethical Framework for a Good AI Society: Opportunities, Risks, Principles, and Recommendations.” *Minds and Machines* 28 (4): 689–707.
- Friedman, Batya, and Helen Nissenbaum. 1996. “Bias in Computer Systems.” *ACM Trans. Inf. Syst.* 14: 330–47. <https://doi.org/10.1145/230538.230561>.
- Ghani, Rayid. 2016. “You Say You Want Transparency and Interpretability?” *Rayid Ghani* (blog). April 29, 2016. <http://www.rayidghani.com/you-say-you-want-transparency-and-interpretability>.
- Goel, Ashok, Brian Creeden, Mithun Kumble, Shanu Salunke, Abhinaya Shetty, and Bryan Wiltgen. 2015. “Using Watson for Enhancing Human-Computer Co-Creativity.” In *2015 AAAI Fall Symposium Series*.
- Gregor, Shirley, and Izak Benbasat. 1999. “Explanations From Intelligent Systems: Theoretical Foundations and Implications for Practice.” *MIS Quarterly* 23 (December): 497–530. <https://doi.org/10.2307/249487>.
- Hager, Gregory D., Ann Drobnis, Fei Fang, Rayid Ghani, Amy Greenwald, Terah Lyons, David C Parkes, et al. 2017. “Artificial Intelligence for Social Good,” 24–24.
- Haque, Albert, Michelle Guo, Alexandre Alahi, Serena Yeung, Zelun Luo, Alisha Rege, Jeffrey Jopling, et al. 2017. “Towards Vision-Based Smart Hospitals: A System for Tracking and Monitoring Hand Hygiene Compliance,” August. <https://arxiv.org/abs/1708.00163v3>.

- Henry, Katharine E., David N. Hager, Peter J. Pronovost, and Suchi Saria. 2015. "A Targeted Real-Time Early Warning Score (TREWScore) for Septic Shock." *Science Translational Medicine* 7 (299): 299ra122-299ra122. <https://doi.org/10.1126/scitranslmed.aab3719>.
- Herlocker, Jonathan L., Joseph A. Konstan, and John Riedl. 2000. "Explaining Collaborative Filtering Recommendations." In *Proceedings of the 2000 ACM Conference on Computer Supported Cooperative Work*, 241–250. ACM.
- Kaye, Jane, Edgar A. Whitley, David Lund, Michael Morrison, Harriet Teare, and Karen Melham. 2015. "Dynamic Consent: A Patient Interface for Twenty-First Century Research Networks." *European Journal of Human Genetics* 23 (2): 141–46. <https://doi.org/10.1038/ejhg.2014.71>.
- King, Thomas C., Nikita Aggarwal, Mariarosaria Taddeo, and Luciano Floridi. 2019. "Artificial Intelligence Crime: An Interdisciplinary Analysis of Foreseeable Threats and Solutions." *Science and Engineering Ethics*, February. <https://doi.org/10.1007/s11948-018-00081-0>.
- Lakkaraju, Himabindu, Everaldo Aguiar, Carl Shan, David Miller, Nasir Bhanpuri, Rayid Ghani, and Kecia L. Addison. 2015. "A Machine Learning Framework to Identify Students at Risk of Adverse Academic Outcomes." In *Proceedings of the 21st ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1909–1918. ACM.
- Lu, Haonan, Mubarak Arshad, Andrew Thornton, Giacomo Avesani, Paula Cunnea, Ed Curry, Fahdi Kanavati, et al. 2019. "A Mathematical-Descriptor of Tumor-Mesosopic-Structure from Computed-Tomography Images Annotates Prognostic- and Molecular-Phenotypes of Epithelial Ovarian Cancer." *Nature Communications* 10 (1): 764. <https://doi.org/10.1038/s41467-019-08718-9>.
- Lum, Kristian, and William Isaac. 2016. "To Predict and Serve?" *Significance* 13 (5): 14–19. <https://doi.org/10.1111/j.1740-9713.2016.00960.x>.
- Lynskey, Orla. 2015. *The Foundations of EU Data Protection Law*. Oxford Studies in European Law. Oxford, New York: Oxford University Press.
- Martinez-Miranda, Juan, and Arantza Aldea. 2005. "Emotions in Human and Artificial Intelligence." *Computers in Human Behavior* 21 (2): 323–41. <https://doi.org/10.1016/j.chb.2004.02.010>.
- McFarlane, Daniel. 1999. "Interruption of People in Human-Computer Interaction: A General Unifying Definition of Human Interruption and Taxonomy," August.
- McFarlane, Daniel, and Kara Latorella. 2002. "The Scope and Importance of Human Interruption in Human-Computer Interaction Design." *Human-Computer Interaction* 17 (March): 1–61. https://doi.org/10.1207/S15327051HCI1701_1.
- Mohanty, Suchitra, and Rahul Bhatia. 2017. "Indian Court's Privacy Ruling Is Blow to Government." *Reuters*, August 25, 2017. <https://www.reuters.com/article/us-india-court-privacy-idUSKCN1B40CE>.
- Neff, Gina, and Peter Nagy. 2016. "Talking to Bots: Symbiotic Agency and the Case of Tay." *International Journal of Communication* 10 (October): 4915–31.
- Nijhawan, Lokesh P, Manthan Janodia, Muddu Krishna, Kishore Bhat, Laxminarayana Bairy, Nayanabhirama Udupa, and Prashant Musmade. 2013. *Informed Consent: Issues and Challenges*. Vol. 4. <https://doi.org/10.4103/2231-4040.116779>.
- Nissenbaum, Helen. 2009. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford University Press.
- . 2011. "A Contextual Approach to Privacy Online." *Daedalus* 140 (4): 32–48.
- Pagallo, Ugo. 2017. "From Automation to Autonomous Systems: A Legal Phenomenology with Problems of Accountability." In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI-17)*, 17–23.
- Pedreshi, Dino, Salvatore Ruggieri, and Franco Turini. 2008. "Discrimination-Aware Data Mining." In , 560–68. ACM. <https://doi.org/10.1145/1401890.1401959>.

- “Pregnancy Mortality Surveillance System | Maternal and Infant Health | CDC.” 2019. January 16, 2019. <https://www.cdc.gov/reproductivehealth/maternalinfanthealth/pregnancy-mortality-surveillance-system.htm>.
- Price, W. Nicholson, and I. Glenn Cohen. 2019. “Privacy in the Age of Medical Big Data.” *Nature Medicine* 25 (1): 37. <https://doi.org/10.1038/s41591-018-0272-7>.
- R. Kerr, Ian. 2003. “Bots, Babes and the Californication of Commerce.” *University of Ottawa Law and Technology Journal* 1 (January).
- Reed, Chris. 2018. “How Should We Regulate Artificial Intelligence?” *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 376 (2128): 20170360.
- Ross, Casey, and Ike Swetlitz. 2017. “IBM Pitched Watson as a Revolution in Cancer Care. It’s Nowhere Close.” STAT. September 5, 2017. <https://www.statnews.com/2017/09/05/watson-ibm-cancer/>.
- “Royal Free - Google DeepMind Trial Failed to Comply with Data Protection Law.” 2017. Information Commissioner’s Office. July 3, 2017. <https://ico.org.uk/about-the-ico/news-and-events/news-and-blogs/2017/07/royal-free-google-deepmind-trial-failed-to-comply-with-data-protection-law/>.
- Shortliffe, Edward H., and Bruce G. Buchanan. 1975. “A Model of Inexact Reasoning in Medicine.” *Mathematical Biosciences* 23 (3): 351–79. [https://doi.org/10.1016/0025-5564\(75\)90047-4](https://doi.org/10.1016/0025-5564(75)90047-4).
- Solove, Daniel J. 2008. *Understanding Privacy*. Vol. 173. Harvard university press Cambridge, MA.
- Strickland, Eliza. 2019. “How IBM Watson Overpromised and Underdelivered on AI Health Care.” IEEE Spectrum: Technology, Engineering, and Science News. April 2, 2019. <https://spectrum.ieee.org/biomedical/diagnostics/how-ibm-watson-overpromised-and-underdelivered-on-ai-health-care>.
- Swearingen, Kirsten, and Rashmi Sinha. 2002. “Interaction Design for Recommender Systems.” In *Designing Interactive Systems*, 6:312–334.
- Tabuchi, Hiroko, and David Gelles. 2019. “Doomed Boeing Jets Lacked 2 Safety Features That Company Sold Only as Extras.” *The New York Times*, April 5, 2019, sec. Business. <https://www.nytimes.com/2019/03/21/business/boeing-safety-features-charge.html>.
- Taddeo, Mariarosaria. 2014. “The Struggle Between Liberties and Authorities in the Information Age.” *Science and Engineering Ethics*, September, 1–14. <https://doi.org/10.1007/s11948-014-9586-0>.
- . 2017. “Trusting Digital Technologies Correctly.” *Minds and Machines* 27 (4): 565–568.
- Taddeo, Mariarosaria, and Luciano Floridi. 2011. “The Case for E-Trust.” *Ethics and Information Technology* 13 (1): 1–3.
- . 2015. “The Debate on the Moral Responsibilities of Online Service Providers.” *Science and Engineering Ethics*, November. <https://doi.org/10.1007/s11948-015-9734-1>.
- . 2018a. “How AI Can Be a Force for Good.” *Science* 361 (6404): 751–752.
- . 2018b. “Regulate Artificial Intelligence to Avert Cyber Arms Race.” *Nature* 556 (7701): 296. <https://doi.org/10.1038/d41586-018-04602-6>.
- Taylor, Linnet, and Dennis Broeders. 2015. “In the Name of Development: Power, Profit and the Datafication of the Global South.” *Geoforum* 64: 229–237.
- The Economist*. 2014. “Waiting on Hold - Ebola and Big Data,” October 27, 2014. <https://www.economist.com/science-and-technology/2014/10/27/waiting-on-hold>.
- Thelisson, Eva, Kirtan Padh, and L. Elisa Celis. 2017. “Regulatory Mechanisms and Algorithms towards Trust in AI/ML.” In *Proceedings of the IJCAI 2017 Workshop on Explainable Artificial Intelligence (XAI)*, Melbourne, Australia.
- Wachter, Sandra, Brent Mittelstadt, and Luciano Floridi. 2016. “Why a Right to Explanation of Automated Decision-Making Does Not Exist in the General Data Protection Regulation.” SSRN Scholarly Paper ID 2903469. Rochester, NY: Social Science Research Network. <https://papers.ssrn.com/abstract=2903469>.

- . 2017. “Why a Right to Explanation of Automated Decision-Making Does Not Exist in the General Data Protection Regulation.” *International Data Privacy Law* 7 (2): 76–99.
- Wang, Yilun, and Michal Kosinski. 2018. “Deep Neural Networks Are More Accurate than Humans at Detecting Sexual Orientation from Facial Images.” *Journal of Personality and Social Psychology* 114 (2): 246.
- Watson, David S., Jenny Krutzinna, Ian N. Bruce, Christopher EM Griffiths, Iain B. McInnes, Michael R. Barnes, and Luciano Floridi. 2019. “Clinical Applications of Machine Learning Algorithms: Beyond the Black Box.” *BMJ* 364 (March): 1886. <https://doi.org/10.1136/bmj.1886>.
- White, Geoff. 2018. “Child Advice Chatbots Fail Sex Abuse Test,” December 11, 2018, sec. Technology. <https://www.bbc.com/news/technology-46507900>.
- Yadav, Amulya, Hau Chan, Albert Jiang, Eric Rice, Ece Kamar, Barbara Grosz, and Milind Tambe. 2016. “POMDPs for Assisting Homeless Shelters – Computational and Deployment Challenges.” In *Autonomous Agents and Multiagent Systems*, edited by Nardine Osman and Carles Sierra, 67–87. Lecture Notes in Computer Science. Springer International Publishing.
- Yadav, Amulya, Hau Chan, Albert Xin Jiang, Haifeng Xu, Eric Rice, and Milind Tambe. 2016. “Using Social Networks to Aid Homeless Shelters: Dynamic Influence Maximization under Uncertainty.” In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, 740–748. International Foundation for Autonomous Agents and Multiagent Systems.
- Yadav, Amulya, Bryan Wilder, Eric Rice, Robin Petering, Jaih Craddock, Amanda Yoshioka-Maxwell, Mary Hemler, Laura Onasch-Vera, Milind Tambe, and Darlene Woo. 2018. “Bridging the Gap Between Theory and Practice in Influence Maximization: Raising Awareness about HIV among Homeless Youth.” In *IJCAI*, 5399–5403.
- Yang, Guang-Zhong, Jim Bellingham, Pierre E. Dupont, Peer Fischer, Luciano Floridi, Robert Full, Neil Jacobstein, et al. 2018. “The Grand Challenges of Science Robotics.” *Science Robotics* 3 (14): eaar7650. <https://doi.org/10.1126/scirobotics.aar7650>.
- Zhou, Wei, and Gaurav Kapoor. 2011. “Detecting Evolutionary Financial Statement Fraud.” *Decision Support Systems*, On quantitative methods for detection of financial fraud, 50 (3): 570–75. <https://doi.org/10.1016/j.dss.2010.08.007>.

APPENDIX: REPRESENTATIVE AI4SG EXAMPLES

Name	Reference	Areas	Relevant factor
1. Field Optimization of the Protection Assistant for Wildlife Security.	(Fang et al. 2016)	Environmental sustainability	1), 4)
2. Identifying Students at Risk of Adverse Academic Outcomes	(Lakkaraju et al. 2015)	Education	4)
3. Health Information for Homeless Youth to Reduce the Spread of HIV	(Yadav, Chan, Xin Jiang, et al. 2016; Yadav et al. 2018)	Poverty, public welfare, public health	4)
4. Interactive activity recognition and prompting to assist people with cognitive disabilities	(Chu et al. 2012)	Disability, public health	3), 4), 7)
5. Virtual teaching assistant experiment	(Eicher, Polepeddi, and Goel 2017)	Education	4), 6)
6. Detecting evolutionary financial statement fraud	(Zhou and Kapoor 2011)	Finance, crime	2)
7. Tracking and monitoring hand hygiene compliance	(Haque et al. 2017)	Health	5)