

AI for the Common Good?! Pitfalls, challenges, and Ethics Pen-Testing*

Bettina Berendt

KU Leuven, Department of Computer Science, Celestijnenlaan 200A, 3001 Heverlee, Belgium
E-mail: bettina.berendt@cs.kuleuven.be

Abstract

Recently, many AI researchers and practitioners have embarked on research visions that involve doing AI for “Good”. This is part of a general drive towards infusing AI research and practice with ethical thinking. One frequent theme in current ethical guidelines is the requirement that AI be good for all, or: contribute to the Common Good. But what is the Common Good, and is it enough to want to be good? Via four lead questions, I will illustrate challenges and pitfalls when determining, from an AI point of view, what the Common Good is and how it can be enhanced by AI. The questions are: What is the problem / What is a problem?, Who defines the problem?, What is the role of knowledge?, and What are important side effects and dynamics? The illustration will use an example from the domain of “AI for Social Good”, more specifically “Data Science for Social Good”. Even if the importance of these questions may be known at an abstract level, they do not get asked sufficiently in practice, as shown by an exploratory study of 99 contributions to recent conferences in the field. Turning these challenges and pitfalls into a positive recommendation, as a conclusion I will draw on another characteristic of computer-science thinking and practice to make these impediments visible and attenuate them: “attacks” as a method for improving design. This results in the proposal of *ethics pen-testing* as a method for helping AI

designs to better contribute to the Common Good.

Keywords: artificial intelligence, machine learning, data science, AI ethics, ethics and ethics codes, risks and impacts, Common Good, AI for Social Good

1 Introduction

Artificial Intelligence (AI) is currently experiencing another “summer” in terms of perceived promises and economic growth. At the same time, there are widespread debates around AI’s perceived risks and negative impacts. In response to the latter, AI researchers and practitioners are paying increasing attention to existing ethics codes, and they are drafting new ones. In addition, many have embarked on research programs that explore how to do AI “for Good”. These two reactions are linked, at a high level, by the understanding that the goal of ethics codes is to encourage and ensure “ethical” professional conduct in the sense of this conduct being “morally good or correct” and “avoiding activities [...] that do harm to people or the environment”.^{1 2}

In addition to the goal to do “good”, many current ethics codes and discussions go further and require that AI contribute to the Common Good. This term is not uniquely (and in many publications not at all) defined, but can be understood as the aim to be good for all.

The purpose of the current article is to investigate more closely the notion of AI for the Common Good by drawing on a wider literature, and to start a deeper discussion in the AI community about this

*To appear in *Paladyn. Journal of Behavioral Robotics*; accepted on 27-10-2018

goal and the way towards it. Towards this purpose, I invite researchers and practitioners to ask four reflective questions of their research practices and projects. These questions can be used as *provocations*: interruptions of the flow of everyday practices designed to “initiate critical reflection [...] on issues that are often otherwise overlooked, obscured or accepted as naturalised practice” [1, p. 225], see [2] for the use of provocations to encourage reflection on big data.

The article is structured as follows.

The Common Good is a notion predating AI. I will start from the definitions given in various AI ethics codes of the Common Good and related notions, and draw on selected discussions in political philosophy for deriving questions about these definitions and their operationalization for AI. These definitions and questions are the subject of Section 2.

Section 3 will provide definitions of other key terms used in the article, including AI, data science, and knowledge. The very general term “AI” will be used to denote research and projects that involve the processing and analysis of knowledge and data, often with machine learning / data mining methods. This interpretation corresponds to the strong representation of data science projects at least in the “AI for Social Good” literature, see Section 5.2. Specific references to data science and machine learning / data mining will be made when necessary.

Contributing to the Common Good is an ambitious and noble aim, and I am convinced that it inspires many researchers and practitioners to act in responsible ways. However, as I will argue in this paper, even with the best of intentions, certain characteristics of AI thinking and practice, coupled with the inherent need to act in politically charged environments, may impede ‘design for the Common Good’. To explain why, Section 4 will detail four specific characteristics, summarized into four lead questions: the problem-solving and solutionism mindset of the engineer, the difficulties of integrating different stakeholders, the role of knowledge, and side effects and dynamics.

Section 5 will validate the importance of the four lead questions via an exploratory survey of 99 contributions to recent conferences on AI and Data Science “for Social Good” or “for Good”, the notions that are currently most similar to AI for the Common Good

and that are sufficiently established to have formed conferences.

Turning these challenges into a positive recommendation, the concluding Section 6 will draw on another characteristic of computer-science thinking and practice to make the impediments visible and attenuate them: “attacks” as a method for improving design. In analogy with penetration attacks, I will propose *ethics pen-testing* as a method for helping AI designs to better contribute to the Common Good. Further, I will argue why the arguments put forward here are characteristic of and relevant for AI and for the goal of enhancing the Common Good, but not restricted to the field or the goal.

2 What is the Common Good?

2.1 The Common Good as a goal for AI

The ambition to be good for all (or at least many) people has become prominent throughout computer science in general and AI in particular. Some examples can be found in ethics codes:

- ACM Code of Ethics and Professional Conduct [3]: “1.1 Contribute to society and human well-being. This principle concerning the quality of life of all people affirms an obligation to protect fundamental human rights and to respect the diversity of all cultures.”
- Asilomar Principles [4]:
 - “23) Common Good: Superintelligence should only be developed in the service of widely shared ethical ideals, and for the benefit of all humanity rather than one state or organization.”
 - “14) Shared Benefit: AI technologies should benefit and empower as many people as possible.”
 - “15) Shared Prosperity: The economic prosperity created by AI should be shared broadly, to benefit all of humanity.”

- Similar ideas are implicit in IEEE Ethically Aligned Design [5, p. 5]: the goal to “develop successful autonomous intelligent systems that will benefit society” and the second General Principle, to “Prioritize the maximum benefit to humanity and the natural environment.”

The first thing to note in these different principles is how differently collectives are referred to. They range from ‘not all of the benefits should accrue to giant internet companies’ (“rather than one state or organization”) to literally “all people” or “all humanity”. “As many people as possible” lies between these extremes, but is underspecified when one does not know what constitutes the possible. Further underspecified terms are the “broadly shared prosperity” and the “widely shared ethical ideals” (see Section 3 for possible referents). The wordings also leave room for different distributions of the benefits, and they make no statements about how to negotiate multiple and possibly conflicting ideals, values, and notions of what is good. Many of these questions have been and are being debated in the wider literature on the Common Good, which is the subject of the following section.

2.2 Some questions regarding the Common Good, inspired by the notion from political philosophy

The Common Good has been discussed widely and controversially by many authors in political philosophy, and it is impossible to survey this literature in the scope of this article.³ Instead, I will very briefly present some issues that raise relevant questions for the interpretation of the concepts proposed in AI.

The Common Good has been defined as “that which benefits society as a whole” [6]. But how are these elements (the “that”, “benefit”, “society”) defined?

Hussain [7] gives more details about the *that*: “the common good is [...] part of an encompassing model for practical reasoning among the members of a political community. [...] The relevant [interests and facilities that serve these interests] together constitute the common good and serve as a shared standpoint

for political deliberation. [...] The relevant facilities may be part of the natural environment (e.g., the atmosphere, a freshwater aquifer, etc.) or human artifact (e.g., hospitals, schools, etc.). But the most important facilities [...] are social institutions and practices.” One example of such institutions and practices is a scheme of private property. Fundamental rights / human *rights* (“basic rights and freedoms”) are parts of the Common Good [7]. Finally, I will use *values* interchangeably with “interests” for the purposes of the present article.

The notion of *benefit* also invites different readings: is it an individual’s or a group’s utility in a welfare consequentialist sense, and/or is it based on values beyond this? (Hussain [7] favors the latter reading, but also reports on alternative conceptualizations of the Common Good.) Finally, the questions of what the boundaries of the relevant *society* (or: political community and its members) are, and of whether to take a welfare consequentialist or other standpoint, and whether and how to account for collective above individual interests [8], tend to receive less attention than others in “for Good” initiatives, and will therefore not be considered further here.

But even the definitional elements of facilities, interests, and practical reasoning raise further questions. The following is a selection that contributed to the choice of lead questions proposed below.

A first question is: Who defines the Common Good (or the interests and facilities) and how? Political philosophy distinguishes between *substantive* and *proceduralist* conceptions of the Common Good. Substantive conceptions specify what factors, goods, values, etc. are beneficial and shared. Proceduralist conceptions instead focus on what procedures are adequate to collectively negotiate and define what is beneficial.

The expression “substantive value” is intended to denote the unassailable status of the value as something that can stand on its own and requires no justification. Yet that status is logically dependent on the attribution of the speaker, who categorizes the value as such. Any such self-supporting value is easily challenged by denying the attribution. Substantive values and their attribution have come under specific political and philosophical attacks after the atrocities

of 20th century authoritarian regimes, who all professed to act in the interest of some common good, an “attempt to make heaven on earth” that “invariably produces hell.” [9]

Proceduralist notions of the Common Good rely on democratic structures and deliberation; it need not be known a priori which facilities and interests will be agreed upon through these processes, see [10]. Even if the focus of proceduralist notions is on process, this does not mean that there are no substantive elements, e.g. [8]. The need for substantive elements can arise from what Popper called the tolerance paradox (if a society is tolerant without limit, this tolerance can be abused or even destroyed by the intolerant). Countermeasures include constraints on the forms the deliberation can take (e.g., that citizens recognize each other as equal and use only reasons that can be accepted by all others [11]) and legal constructs that enable and require a country’s political bodies to protect the political order against those who want to abolish them, such as constitutional clauses that cannot be abolished even by a majority (“militant democracy”, cf. [12]).

Another distinction is that between *communal* and *distributive* conceptions of the Common Good. A communal conception takes the Common Good interests to be interests that citizens have as citizens, whereas a distributive conception is based on the acknowledgement that citizens belong to various groups with distinct interests, that these interests compete for the facilities and resources and may pose different demands, and decisions and allocations need to be made according to some distributive principle [7].

2.2.1 From questions about the Common Good to questions about AI for the Common Good

The philosophical considerations about the Common Good that have been summarized very briefly in the previous section served as starting points for the questions proposed in Section 4. Here, I will give an overview of the link between the considerations and questions.

The considerations above indicate that purely substantive accounts of the Common Good are problem-

atic, that procedures are important, and that groups with different interests and demands may have different notions of the Common Good. These groups correspond to what computer science calls stakeholders. These considerations were one inspiration for the first two lead questions: What is the problem, and who defines it (see Sections 4.1 and 4.2)? A second inspiration was that these same two questions have proved constructive in interdisciplinary collaboration around a specific Common Good interest and facility: “privacy” [13].

Note that the definitional duality of “Good” and “problem” introduced in the previous paragraph is frequent in AI: some value or aspect of the Common Good is missing, deficient, or under attack, and this constitutes a problem. The problem then prompts a search for a technological contribution to solving or at least addressing this problem. The focus on a (usually technological) solution is the reason to ask a modified version of the first lead question: What is a problem in the first place (see Section 4.3)?

“AI for the Common Good” is understood here (and in the surveyed literature) in an engineering sense. Thus, AI methods, technology, and their deployment cannot be an interest, but a facility (or part of it) that serves an interest. This raises the question: what kind of facility is or should this be? I will argue that today, this is mostly some form of knowledge that is then fed into further decision processes. (Another candidate are autonomous systems, which would require a different analysis.) This inspired the third lead question concerning what the role of knowledge is (see Section 4.4).

The fourth lead question asks about important side effects and dynamics (see Section 4.5). This can be related to the Common Good literature in that this literature also investigates what is likely to happen under different structures of people acting, deliberating, deciding, and collaborating.

2.3 AI and Data Science for (Social) Good

At this point in the argument, one would normally investigate how concrete projects or initiatives (rather than abstract ethics codes) define the Common Good.

However, the call to develop and deploy AI for the Common Good has, to the best of my knowledge, not yet led to research programs or publications under that name. At the time of writing of this article (August 2018), a Google search returned three texts. The phrasing “AI for Common Good” has been used in the titles of two recent white papers, one prepared for attendees of the 2018 World Economic Forum Annual Meeting [14], and one by North Highland Consulting [15]. Both focus on highlighting threats posed by Artificial Intelligence; what the Common Good is and how to use AI towards it is not explicated. In addition, an entry in the *Communications of the ACM*’s news, titled “AI for the Common Good” [16], reports on the AI for Good Global Summit, whose goal definition is given below.

However, related ideas have a longer tradition and have led to several conferences and research programs. To outline the field, I have selected four that I believe are most influential and representative of “AI for (some version of) Good”. The selection was based on the duration of the initiative (at least two editions) and/or the backing by an important professional association (AAAI) or an important international actor (the UN). Since the initiatives present their definitions on their Web pages rather than in scientific publications, some interpretation is needed and will be supplied in the following paragraphs.

Initiatives around AI and Data Science will be presented, for several reasons. First, “for Social Good” originated as an initiative from data science, second, data science is one key area of, or related to, current AI (for details, see the definitions in Section 3), such that, third, many contributions to conferences on AI for (Social) Good are or contain data science.

The *Data Science for Social Good* (DSSG) initiative has organized, since 2013, an annual “summer program for aspiring data scientists to work on data mining, machine learning, big data, and data science projects with social impact. Working closely with governments and nonprofits, fellows take on real-world problems in education, health, energy, transportation, and more.”⁴ The first part of this definition is strictly speaking not very specific, since many uses of AI have social impact, including large social networks, search engines, and (inter)national surveil-

lance programs. This shifts the definitional core of “Social Good” to the domains of projects (such as education, health, energy and transportation) and to the actors (governments and nonprofits) who are likely to be the ones to define project goals and/or who control and provide the data.

A definition via *domains* and *actors* is also used by the Data Science for Social Good (SoGood) workshop series that has so far seen three consecutive editions at ECML PKDD, a major conference on machine learning, data mining, and data science: “how Data Science can and does contribute to social good in its widest sense, including areas such as: Public safety and disaster relief, Access to food, water, and utilities, Efficiency and sustainability, Government transparency, Data journalism, Economic development, Education, Social services, Healthcare. We are interested both in non-profit projects and in projects that, that while not defined as non-profit, still have Social Good as their main focus, and so have managed to build a sustainable business model.”⁵

A focus on DSSG projects’ *problem-solving* is suggested by [17]: DSSG consists of “attempts to solve complex social problems through the use of increasingly available, increasingly combinable, and increasingly computable digital data”.

With a method scope of AI in general (rather than DS in particular), the Association for the Advancement of Artificial Intelligence held a spring symposium on “AI for the Social Good” in 2017.⁶ The AAAI Spring Symposia center on emerging topics in AI; hence, this is an indication of the endorsement of the field, by a major professional association. “AI for the Social Good” is defined as AI “addressing societal challenges, which have not yet received significant attention by the AI community or by the constellation of AI sub-communities, [the use of] AI methods to tackle unsolved societal challenges in a measurable manner.”⁷ Another venue defines the field by declaring “almost any real-world problem, which is important for society’s benefit, and could potentially be solved using AI techniques, [to be] within the ambit of this symposium.”⁸ This definition reiterates the idea of “benefit for society”, see Section 2.2, and the focus on problem-solving.

With a method scope of “Good” in general (rather

than “Social Good” in particular), the ITU, the UN agency responsible for issues that concern information and communication technologies, leads the “AI for Global Good Summit”. The Summit has so far been organized twice (2017 and 2018). The goal is described as “AI innovation [being] central to the achievement of the United Nations’ Sustainable Development Goals (SDGs) by capitalizing on the unprecedented quantities of data now being generated on sentiment behavior, human health, commerce, communications, migration and more”, including goals such as “no poverty”, “zero hunger”, and “good health and well-being”.⁹ These and most of the 14 other SDG goals have a *substantive* focus.

More specific societal goals, for example fairness (non-discrimination), are pursued by research communities such as Fairness, Accountability and Transparency in Machine Learning and beyond.¹⁰ Another example is the protection of privacy as the goal of various research communities including (in DS) privacy-preserving data mining and data publishing. In addition, funding programs with similar goals exist. I have been part of two “projects with a primary societal finality” funded by the Flemish Science Council FWO. In these projects, multidisciplinary consortia (of which AI was only one partner) investigated privacy in online social networks and diversity in media, respectively. The methodological and ethical debates in these projects have been an important source of inspiration for the current article. Since these goals are quite distinct from (or at least much more specific than) the Common Good, these specific notions of the Good will not be investigated further here.

In sum, “the Common Good” is referred to as a goal for AI in current publications, but not defined. Concrete current initiatives refer to the “Social Good” or simply “the good”, circumscribing it via problem domains and the identity of the non-academic project partners (nonprofits or governments), or via substantive goals agreed upon at UN level. It appears that, as a common denominator, the intended beneficiaries of AI for Good, for Social Good, etc. will generally not be the ones who directly pay for the development or use of this AI. Thus, unlike for example in commercial application areas, no market price can serve as indicator of value. “The

Social Good” is then the indicator of such value.

Healthcare is an interesting example: while it can be a profit-oriented business model, the focus in the domain of AI for (Social) Good appears to lie on the provision of healthcare to broader sections of society (see “for all” above). The historical experience suggests that such a provision requires some kind of national insurance financing scheme based on solidarity rather than payment-for-service. Thus, is any contribution of AI to better healthcare methods already “AI for Good”, or is more needed? It is likely that such questions will be asked in the further development of the field.

3 Terminology: AI, Data Science, knowledge, and ethics-in-AI

In this section, further key terms used in this article will be defined.

Artificial Intelligence (AI) is “the ability of a digital computer or computer-controlled robot to perform tasks commonly associated with intelligent beings. The term is frequently applied to the project of developing systems endowed with the intellectual processes characteristic of humans, such as the ability to reason, discover meaning, generalize, or learn from past experience” [18]. In line with the ACM Computer-Science subjects classification, I regard AI as a subfield of computer science, and as such a field at the intersection of science and engineering. *Machine Learning* (involved in particular in the last three characteristics of the preceding list) is a field of AI, as encoded for example in the ACM Computer-Science subjects classification in its most recent (2012) version.¹¹

Data Science (DS) is understood here as a subfield of AI, or more specifically as a field that (a) is situated in many universities within AI groups and (b) draws heavily on methods developed or used in machine learning and data mining. (The machine learning aspect corresponds to the focus on learning models of data, and the data mining aspect to the focus on entire knowledge-discovery workflows.) More gen-

erally, data science has been defined as “the science (or study) of data” and “a new interdisciplinary field that synthesizes and builds on statistics, informatics, computing, communication, management, and sociology to study data and its environments (including domains and other contextual aspects, such as organizational and social aspects) in order to transform data to insights and decisions by following a data-to-knowledge-to-wisdom thinking and methodology” [19]. Conway [20], in a frequently cited online source, described data science by means of a Venn diagram: with regard to machine learning, data science is situated in the intersection of machine learning and substantive expertise.

Knowledge is understood in two ways. On the one hand, I will refer to a notion of knowledge as used in psychology: “a structured collection of information that can be acquired through learning, perception or reasoning” [21], understood to be held by a human agent (mental representation). It is the knowledge about something in the world, the expertise [17] that generally draws on many sources and fields, such as different academic disciplines. Such human knowledge can be both *input* to a research or development activity, and its eventual *result*.

On the other hand, I will refer to knowledge as the more immediate *output* of an AI or DS activity. Russell and Norvig [22, p. 16] implicitly define knowledge as a structured collection of “information [...] put into a form that a computer can reason with”. The field of knowledge discovery from databases and the related fields of data mining and data science focus on knowledge in the sense of “novel, valid, potentially useful, and ultimately understandable patterns in data” [23] – where the intended recipient, who can use and understand these patterns as structured representations, is often but not necessarily human. In all these meanings, knowledge is structured information, useful and/or understandable to a person or machine.

In this article, the very general term “AI” will be used to denote research and projects that involve the processing and analysis of knowledge and data, often with machine learning / data mining methods. This interpretation corresponds to the strong representation of data science projects at least in the “AI for

Social Good” literature, see Section 5.2. Specific references to data science and machine learning / data mining will be made when necessary.

A final note concerns the question: *Whose ethics?*

Robotics as such is not in the focus of the present article, but robots are relevant to the focus on AI and ethics. Not all artificial intelligence is incorporated into robotics, and not all robots are artificially intelligent. However, the intersection is large and relevant, and I regard such AI robots as typical representatives of what the Asilomar Principles call “highly autonomous AI systems”. Regarding these, three types of ethics are relevant.

The first type is the professional ethics of the researcher or practitioner (often referred to as computer ethics). This is guided by principles such as the Asilomar Principle 11) “Human Values: AI systems should be designed and operated so as to be compatible with ideals of human dignity, rights, freedoms, and cultural diversity.” Arguably, these ideals are widely shared, codified for example in the Universal Declaration of Human Rights.¹² Thus, one may interpret the “widely shared ethical ideals” of Principle 23) as consisting of these four, and possibly also others. (According to Principle 23), “[s]uperintelligence should only be developed in the service of widely shared ethical ideals and for the benefit of all humanity rather than one state or organization.”)

The second type is the professional ethics of a researcher or practitioner who develops robots. Following Veruggio [24], I refer to this as roboethics: “Roboethics is not the ethics of robots nor any artificial ethics, but it is the human ethics of the robots’ designers, manufacturers, and users.” Asilomar Principle 10) constrains these design, manufacturing and use activities by positing “Value Alignment: Highly autonomous AI systems should be designed so that their goals and behaviors can be assured to align with human values throughout their operation.”

Note that this formulation focuses on the goals and behaviors of a robot and asks that these be aligned with human values (presumably those of Principle 11) listed above), and that it does not make a commitment as to whether these goals and behaviors stem from the humans designing, manufacturing, or using the robot, or from the robot itself. Thus, the ques-

tion of whether machine ethics [25, 26, 27] as a third type of ethics, the ethics of a robot (or in fact any AI system), exists and if so, what its properties are, is left open. In line with this, the remainder of this article will focus on computer ethics / roboethics as human ethics in the sense described above.

4 How to create “AI for the Common Good”: Four lead questions

In this section, I will analyze four specific characteristics of AI thinking and practice that challenge and may impede design for the Common Good: the problem-solving and solutionism mindset of the engineer, the integration of stakeholders, the role of knowledge, and side effects and dynamics.

The questions will be illustrated by references to a running example from the domain of “AI for Social Good”. The example itself is, intentionally, not a real example in the sense of being the contents of a specific AI paper, report or otherwise – because the point of the present article is not to denigrate the merits of any particular project. Instead, the example is a fictitious synopsis of uses of AI/DS in various contexts, with these uses all focusing on the same issue. The example centers on drugs, considered by some to be “public enemy number one”¹³, that is, the ultimate “Common Bad”, whose absence would surely enhance the Common Good. While traditionally, such statements targeted illegal drugs, the recent US opioid crisis, which was declared a Nationwide Public Health Emergency¹⁴, has highlighted how a similar problem can originate from a substance that may be legally prescribed or illegally peddled. The opioid crisis also illustrates how public health and criminal justice issues continue to interact: Within three paragraphs of one political speech, US President Trump lauded an initiative that caused people to turn in more than 900,000 pounds of unused or expired prescription drugs, and the arrest of criminal aliens with 76,000 charges and convictions for dangerous drug crimes.¹⁵

The rhetoric around the opioid crisis has brought

one question back into sharp focus: There is a problem, but what exactly is wrong? This leads to the first lead question Q1.

4.1 Q1: “What is the problem?”

Consider the example: What is “the drug problem”? Here is a non-exhaustive list of candidates:

1. People use drugs.
2. People sell drugs.
3. Certain people (e.g. the poor, black people, ...) use drugs.
4. Drug users commit crimes.
5. Drug users become homeless, ill, ...
6. Drug users die (earlier than they would have without drugs).
7. There aren’t enough drugs available.

These alternative definitions are designed to represent, in a simplified way, the views of different people who are affected by drug usage and its consequences.¹⁶ If, therefore, a system to “solve” the drug problem, by AI or otherwise, is created, the question of who defines the problem counts.

4.2 Q2: Who defines the problem?

A necessary condition for designing systems that further the Common Good is to hear the voices of multiple and diverse people who will be affected by the system. The integration of multiple stakeholders in requirements engineering (for an overview, see [28]) are therefore increasingly required by ethics codes as well as laws. Codes such as the ACM Code of Conduct [3]¹⁷, the AOIR Recommendations for Ethical Decision-Making and Internet Research [29]¹⁸, and the IEEE Ethically Aligned Design guidelines [5]¹⁹ explicitly call for this. A special form of multi-stakeholder requirements elicitation has recently even attained the status of a legal obligation: the European Union’s new data protection law, the General Data Protection Regulation (GDPR), requires

a data protection impact assessment before personal data are collected and processed. (While the current article focuses on ethics codes, I will make some references to the GDPR as an example of a current and wide-ranging attempt to codify rules for technology, including AI, to protect individuals' rights and freedoms.)

As part of such processes, “[i]t is understood that there will be clashes of values and norms when identifying, implementing, and evaluating these systems (a state often referred to as ‘moral overload’)” [5, p. 23], so conflict resolution methods and processes are required. The study of *democratic* methods for gathering and negotiating requirements is a subfield of requirements engineering [30]. However, an ongoing challenge remains: how to best support democratic deliberation and conceptions of distributive justice with software and/or software engineering methods.

Conflict identification and resolution become more difficult when stakeholders are differentially able to cause a clash in the first place, because they are embedded in socio-technical systems differently and differ in their abilities to perceive and voice their values and norms. As an example, consider imprisoned drug users as one of the relevant groups in [31]. Further problems arise when affected communities and individuals are outside the boundaries of the society deemed relevant in the respective notion of the Common Good. For example, inhabitants of countries such as Colombia, in which drugs are grown and who suffer from the local effects of drug cartels' power, have argued that they are outside the consideration of drug consumers in the West [32].

In addition, it is questionable whether talking to different stakeholders is enough, because it may not affect the structural mold into which these different stakeholders' utterances will be put: the very notion of what a problem – any problem – is. This will be investigated next.

4.3 Q1': “What is a problem?” (and thereby: What is the problem – revisited)

While conceivably no AI researcher or practitioner would be as preposterous as claiming to “solve the drug problem”, AI approaches do focus on a version of the problem (usually implicitly specified to be smaller). This is evidenced by the above-mentioned reference to the use of AI on “[a]lmost any real-world problem, which is important for society's benefit, and could potentially be solved using AI techniques”. This is germane to the discipline and the “problem-solving mindset” of the engineer.

4.3.1 Problem-solving

In everyday language, a problem is “a matter or situation regarded as unwelcome or harmful and needing to be dealt with and overcome”.²⁰ In some cases, “dealing with and overcoming” may be relatively straightforward. Using another drugs example: if the problem is a person's breathing being suspended due to an overdose, this problem can be dealt with and overcome by the proper administration of Naloxone. However, most real-world problems, including many around drug overdosing, are more complex. For example, subjecting a heroin addict to a methadone program can “deal with” the heroin addiction, but it does not necessarily “overcome” it. In addition, many real-world problems are open-ended. For example, there is probably no “overcoming” the fact that in any society, many people overuse or misuse legal or illegal drugs – but this does not alleviate societies from the responsibility to “deal with” drug usage.

On the other end of the scale, there are chess problems: “an arrangement of pieces in which the solver has to achieve a specified result” or mathematics/physics problems: “an inquiry starting from given conditions to investigate or demonstrate a fact, result, or law.”

Engineering problems may, but do not have to, start from the everyday notion — but the engineering approach rests on transforming whatever the starting point is into a well-defined designation: moving from

a stateA (undesirable situation) to a stateB (desirable situation) [33]. Only once the “problem” has this well-defined shape, can the engineer begin to “solve” it. This is what makes engineering precise, graspable and powerful.

However, most social problems are not chess problems, and they do not exist in context-free structures. Therefore, en route to the definition of an engineering problem, one must usually make some assumptions that are hard or even impossible to formalize. Ex post, the ambivalence that existed in the beginning tends to be cognitively minimized and the result is taken to be *the* truth (even if it did result from decisions that might just as well have been made otherwise), cf. [34]. The result is, partly, that the definition of the problem often arose from the opinions of only one or only a few stakeholders. Multi-stakeholder methods and participatory software design are approaches for addressing this issue. However, regardless of which and how many stakeholders have been consulted, formalization remains a necessary step. The risk of conceiving of social problems in terms of engineering problems is to blind oneself to the vagaries of the formalization step, and to fail to consider alternatives to the chosen formalization.

4.3.2 Problems and “solutions”

Likewise, what do we consider as “solutions”? This will depend on the context in which problem-solving takes place. Consider several standard problem-solution pairs, with solution approaches from law, law enforcement, and public health, in different countries. The current discussion will focus on illegal drugs. In general, problems 1-7 above exist for legal drugs too, but are addressed differently.

In the War on Drugs that began in the Philippines in 2016, attention has focused on the first and second of the problem versions from Section 4.1 above (henceforth, #1 and #2). The “solution” proposed in the election campaign of President Duterte as well as enacted in a large number of cases was to stifle both supply and demand by killing drug dealers and drug users, cf. [35]. Another “solution” approach consists of criminalization and incarceration laws and policies for drug dealing and use, even of small quantities, as

in the US during recent decades. Some countries exempt the ownership and consumption of small quantities of illegal drugs from prosecution (i.e. prioritize #2 in the problem definition). Various societal actors and authors (e.g. [36, 37]) have argued that problem definition #3 stands behind US laws that penalize the use and dealing of drugs traditionally associated with poor and black users (crack cocaine) far heavier than that of similar drugs traditionally more prevalent among affluent white users (powder cocaine). The identity of drug users in problem definition #3 can even become associated with a (re-)framing both of “the problem” and “the solution”. This point has been made after US President Trump, in October 2017, declared the opioid crisis (which affects many poor white people) a public health emergency rather than another type of drugs on which to wage war [38].

#4 can be a problem in at least two ways: #4.1 when drug users commit crimes such as theft or prostitution to finance their addiction, or #4.2 when intoxicated people become uninhibited and/or aggressive and then commit crimes such as assault and murder. #4.2 is an often-voiced observation concerning the legal drug alcohol, and “solutions” contain penalties for driving when intoxicated. (Beyond that, for legal drugs there tends to be a separation between permitting the intoxication as an expression of personal freedoms, and the sanctioning of crimes if they occur.) #4.1, on the other hand, is in many cases tightly linked with #5, and programs such as substituting methadone for heroin in (a) legal, (b) insurance-covered, and (c) medically administered ways are offered as partial “solutions”. If #6 is considered the main problem, the question for a solution may need to turn from the ubiquitous attempts to reduce consumption to accepting the fact that consumption and overdosing happen, and counteract the lethal effects of overdoses as they appear.

These approaches from law, law enforcement and healthcare are not, or not in their entirety, AI-based. AI is used to address parts of the problems, as when predictive policing software recommends where to patrol for drugs [39]. The AI models underlying such software need well-defined and accessible data and well-defined objective functions whose maximization constitutes a “solution”. This can lead to police de-

ciding to patrol where there have been many drug-related arrests in the past, rather than where there has been much drug usage, and it will commit countermeasures against such biased decisions to one formalization of “fairness”, which may mean that other notions of fairness are violated [40, 41]. Such side-effects of AI “solutions” should be kept in mind. To keep the unavoidable restrictions more visible, references to “problems” and “solutions” should be replaced by terms that are more clearly technical and limited in scope, such as “the task to be done”. In the example, the task could be to decide where to patrol. Once the task is clear, the question becomes what to do or communicate.

4.4 Q3: What is the role of knowledge?

Above, I have argued that the transformation of a social problem into a formal problem poses challenges when the goal is to contribute to the Common Good. In this section, I will study two challenges related to this transformation, both of them related to knowledge. The first concerns the knowledge that enters (or fails to enter) into the transformation into a formal problem. The second concerns the reverse direction: what happens when the output or “solution” of this formal problem is a piece of generated knowledge? And what side-effects arise when both input and output are (necessarily) entangled with the knowledge of the AI developer, funder, ...? How may AI methods themselves affect such knowledge effects? What consequences may these constellations have on the Common Good?

AI and in particular DS are strongly linked to knowledge: The goal of AI is often described in a procedure-oriented way, such as in the definition presented in Section 3: to “[develop] systems endowed with the intellectual processes characteristic of humans, such as the ability to reason, discover meaning, generalize, or learn from past experience”. Yet, early on in the history of AI, it became clear that only focusing on algorithms (e.g. for reasoning) will not produce intelligence, but that strong knowledge bases are also needed. Terminology and valuations of logics-based approaches change over time, and cur-

rently, the reliance on knowledge is often expressed as a reliance on (big) data instead. In other words, AI systems and their outputs are considered useful if they work on large and rich data or knowledge, and if their outputs present new information that humans derive further knowledge from and act upon.²¹

4.4.1 The power of knowledge: AI and framing effects

The framing of a problem denotes the way it is described. The framing of a problem (and its associated “solutions”) has a powerful effect on how people – and thus the public as those who may benefit, or be harmed by an AI system – perceive the world and act in it, e.g., [42]. Frames can themselves be objects of knowledge, that is, statements can be made about them and discussed. For example, different frames can be identified and compared as alternatives. Designers should be aware that their knowledge-based methods operate in an environment filled with politically and otherwise induced frames and that association-based methods tend to reinforce these frames. These frames come loaded with certain versions of the notion of the Common Good (and many more blind spots regarding it). To work in the interest of the Common Good, an AI researcher or practitioner should be aware of this fact and make conscious (and transparent) choices about whether to sustain frames or expose them.

Framing interacts with AI for the Common Good on multiple levels. First, frames contribute – or could contribute – to problem definition. Second, AI, by the way it processes knowledge, can serve to reinforce such frames. Third, frames operate not only on the level of problems and solutions, but also on the level of methods. These effects will be considered in turn.

Frames emphasize certain aspects of reality, and they suppress or even block others. One consequence is that specific views – including but not limited to the perception of what the problem is by specific stakeholder groups – can remain suppressed, or that the knowledge that certain solutions do not work is blocked.

An example is the trope of the “War on Drugs”. Launched in 1969 by then-US president Richard

Nixon, this may have been influenced by concerns over public health and the suffering induced by abuses of illegal drugs. However, an alternative view has co-existed with this, succinctly described by Nixon’s then counsel and Assistant to the President for Domestic Affairs: “The Nixon campaign in 1968, and the Nixon White House after that, had two enemies: the antiwar left and black people. [...] We knew we couldn’t make it illegal to be either against the war or black, but by getting the public to associate the hippies with marijuana and blacks with heroin, and then criminalizing both heavily, we could disrupt those communities. We could arrest their leaders, raid their homes, break up their meetings, and vilify them night after night on the evening news. Did we know we were lying about the drugs? Of course we did.” [43]²². Baum’s view changes the roles of pieces of knowledge: In the War on Drugs frame, the drug-taking is the problem, and the solution is (or at least involves) arresting people, raiding their homes, breaking up their meetings, and using the media to vilify them. In the Baum frame, the existence or growing influence or expectedly growing influence of antiwar and black citizens is the problem, and an (at least intermediate) solution is to wage a War on Drugs. Independently of whether one follows Baum’s sinister interpretation of the political will of the Nixon campaign, there is now a more widespread acceptance of criminalization and incarceration not having “solved” anything at all [43]: (a) legal alternatives to criminalization exist, have been or are being tested in different countries, and have often led to measurable improvements in health and crime statistics, (b) legal drugs (in particular alcohol) and prescription drugs (such as opioid painkillers) affect far more people than illegal drugs and cause enormous human suffering and economic costs, and (c) the “cure” of criminalization imparts suffering and creates new problems, also for the Common Good.

Frames tend to be reiterated and “echoed”, and they can survive even in the face of clear scientific evidence that contradicts a frame. Arguably, much of current drug policy rests on reiterated, “echoed” frames that have been constructed and perpetuated for a number of decades now. One example is the perpetuation of the focus on criminal justice even in

the context of the National *Public Health* Emergency proclaimed by US President Trump, see above. Another example is the controversy around David Nutt’s two studies in *The Lancet* arguing that alcohol and tobacco are far more dangerous than common illegal drugs. Among other things, Nutt was dismissed from the UK government Advisory Council on the Misuse of Drugs, see the summary and links in [44].

Mass media have long been known to be influential in relaying, reinforcing and maintaining frames [45], creating “echo chambers”. Social media have recently been described as intensifying the echo-chamber effects that media often have anyway, e.g. [46, 47, 48]. This is where AI enters: recommender systems (the backbone of modern social media platforms’ approach to addressing information overload) work on associations learned from past data and thereby tend to further strengthen these effects, e.g. [49].

Framing operates not only at the level of problems and solutions, but also on the method level. This can have wide-ranging effects on decisions for example in data-science related projects. For example, structuring a decision-making process and tool by first identifying potential harms at a general level and then weighing them against specific and contextual potential benefits “would always, it seemed, come out in favour of intervening and therefore in favor of the data sharing that would enable intervention” [50, pp. 4-5]. In an analysis of big data approaches to epidemics in high-income vs. low-/middle-income countries, the conception of populations as well-informed individuals versus as pathogen-carrying groups may imply that “big data models built to facilitate individuals’ well-being and autonomy instead would constitute perfect tools for mass control and surveillance” [51, p. 30]. Finally, the correlation-based nature of data mining itself has effects. Data science models used in the criminal justice sector have been criticized widely for their effects of reproducing societal biases against minorities, cf. the proceedings of the FAT(ML) conferences as a specific branch of AI for Social Good, see Section 2.3. As Barabas et al. [52] point out, there are however risks for all, regardless of minority status, when the underlying epistemic assumption is that persons “have” crimi-

nal tendencies innately correlated to their features and the intervention focusses on decisions about bail or incarceration, rather than on causal factors and the possibility to change them via effective diagnosis and intervention of criminogenic needs.

In sum, frames affect AI projects, and they can be reinforced by AI techniques. Framing may be unavoidable, but it needs to be reflected and – if appropriate – counteracted. Methods such as those designed to detect frames and bias, and to increase diversity in recommendation [53], can be components of a more critical stance towards these mechanisms.

4.4.2 Limits of imparting knowledge: does it work?

Given the strong effect of framing, someone who ignores frames and believes that they can immediately discern “the facts” in information they consume, or can immediately convey “the facts” in information they produce and communicate, *underrates* knowledge. At the same time, many presentations of AI tend to *overrate* knowledge in the sense that they suggest that the existence of knowledge can by itself solve problems.

Such overrating can occur when knowledge is imparted for awareness raising. This is the goal of awareness tools in general (see [54] for an overview specifically with regard to privacy awareness tools) and today is found in many quantified-self apps. Basically, everybody (including drug users) knows that drugs are bad for health, family life, socio-economic status, etc. – but this does not stop an addict from consuming the drug when the opportunity arises. In general, the limitations of “just informing” have been studied intensively in recent years by, for example, behavioral economists, and alternatives such as “nudging” have been investigated. These approaches rest on acknowledging that “knowledge is not all” and that decision-making is influenced by a wider range of factors than just classical rationality. One important group of factors are social influences. Many quantified-self apps and related applications, including those around substance abuse or addiction problems²³, try not only or not at all to impart knowledge, but rather to help build and maintain social-support

groups.

In their focus on using IT as communication technology, many of these apps are not AI-based. So where does AI come in? A very good example is Bird et al.’s [31] use of data science methods informed by definition #6 of “the drug problem”. The authors used a data-science analysis to show the prevalence of lethal overdoses among addicts recently released from prison (at which time addicts are even more vulnerable than usual due to the enforced abstinence while in prison), and then instead of arguing for an awareness campaign “to avoid drugs” or “avoid overdosing”, handed out emergency overdose kits and gave basic information on how to administer the antidote. They showed, again with methods from data science, the effectiveness of their intervention: the number of deaths decreased significantly.

All these “solutions” rely on certain definitions or framings of problem and solution, and all of them depend on various factors determining decision making. Data scientists and AI designers can draw on methods for supporting these factors and decisions studied in human-computer interaction [55], but as the Bird et al. example suggests, they should also be ready to think outside the box and embrace solutions in which they may only shine as diligent data analysts in the background, rather than as providers of smart knowledge-based tools in the foreground.

But even when imparting knowledge works, is it always a good thing?

4.4.3 Limits of imparting knowledge: is it good?

AI and in particular DS often appear to operate on the assumption “The more knowledge, the better”. This idea is applied to individuals as potential holders of knowledge, and it also impinges on the idea of the Common Good: “The more knowledge society has, the better”. But is this always the case?

At the individual level, there are certain well-known problems. First and as argued in the previous section, imparting knowledge may not work in the envisaged way. It may also manipulate people and negatively impact their autonomy, or hurt them in other ways [13]. It may place undue burden on them

by making them responsible for tasks they lack the mental, financial, temporal, etc. capabilities for (“responsibilization”, [56]). A growing number of legal and ethics guidelines recognize such limits. These include culturally-grounded restrictions on imparting knowledge [57] as well as “the right to not know” in bioethics. Not knowing certain things is also recognized as a helper against unconscious biases, and it can therefore have economic advantages. The properties of such forms and conventions of not-knowing are investigated in the field of ignorance studies [58].

If imparting knowledge is not necessarily beneficial, a parsimony principle can be useful: focus on the task at hand and get and use the knowledge needed for it, but not more. This principle is inspired by ethical, legal, and general intellectual principles. The Nuremberg Code, an early and highly influential code of research ethics, posits: “The experiment should aim at positive results for society that cannot be procured in some other way.”²⁴ The legal principle of proportionality (which pervades laws in general, and is particularly clearly adaptable to current purposes when a knowledge-based activity interferes with the fundamental right to data protection) says that the measure should be necessary to reach the goal. Similarly, in data protection principles and laws such as the GDPR, data minimization (collecting and using as little personal data as possible for the task at hand) is a guiding principle. Finally, arguably minimalism is considered a scientific virtue, expressed by general principles such as Occam’s razor down to specific topics such as zero-knowledge proofs.

4.5 Q4: What are important side effects and dynamics?

Although computer science arose from cybernetics, the study of systems and feedback loops, much of today’s computer science rests on surprisingly linear and short-term cause-effect relationships. This is probably due to that other basic principle of the natural and engineering sciences: divide and conquer, that is, split problems into parts and address these separately. However, side effects and dynamics of applications are becoming more visible. Again, these may negatively impact the overall effect of AI sys-

tems and thereby reduce, annihilate, or even reverse positive effects on the Common Good.

I will illustrate these with some examples from our example domain.

First, by any version of “the drug problem” relating to people and their behavior, inevitably many personal data will be collected and processed. This raises data protection issues that could well outweigh any positive benefits. As an example, consider social network mining, which is currently a popular method also for studying drug usage, see for example [59]. In the event of a data leak²⁷, such social network mining methods could be used to derive inferences about individuals’ drug-related behaviors or propensities that may damage reputations and affect lives. Alternatively, social-media users could receive targeted advertising based on their supposed propensities, and vulnerabilities could be exploited. Importantly, the question is not so much whether the mining methods return true knowledge, or whether the validity of the targeted advertising can be demonstrated: the attempt at manipulation itself may present the problem. This is a lesson learned from the history of the Cambridge Analytica case, in which ideas from an academic project in which social media were mined to predict personality [60] were later supposedly used for psychometrically micro-targeted election advertising [61].

Second, it is by now well-known that “objective” big data analyses are likely to reproduce the biases in the data they learn from (thus violating the right to non-discrimination) [62]. This has been argued for a wide range of big-data applications [63] and shown with simulations for example for drug patrols [39]: A predictive-policing application learns from past data that arrests have occurred frequently in certain areas, and it proposes that police patrol these areas preferentially. This leads to more arrests in these areas, which in turn feeds the learning to propose to patrol them, etc. In general, since the deployment of big-data analyses will itself create data that then become input to further data analyses, this can easily create vicious-cycle phenomena that have been observed by sociologists for long, dynamics that can perpetuate or even aggravate bias and discrimination [64]. A promising research direction for breaking such feed-

	Engineering problem solving approach	Dealing with social problems
Goal function	well-defined	often involves unresolvable ethical dilemmas, continued re-negotiation
Method	can be black box	requires fairness, transparency, accountability
Decomposition	modular	usually interdependent
Delegation	can be fully delegated	(at least some) participation required
Ethics codes and assessments	canonical, often a burdensome afterthought	starting points ²⁵ , guidelines ²⁶
Solvability	Problems can be solved.	Some problems can only be addressed.

Table 1: Solving engineering problems vs. dealing with social problems

back loops, drawing on reinforcement learning, proposes a different strategy for patrolling that could help to also detect the cases in the (initially) less likely areas [65]. It will be interesting to see how such strategies can be put into practice, and what effects this will have.

Third, self-reinforcing feedback loops can occur not only at the level of data, but also regarding technology use. Problematic drug usage is a form of addictive behavior. But if someone decides to control their drug use through an app, this can contribute to addictive forms of Internet usage, and the question arises whether this substitution is sensible. Vice versa, learning to use an app ecosystem in a way that is conducive to one’s well-being, could help overcome substance abuse. It is an open question which factors contribute to these dynamics playing out in vicious or virtuous feedback loops.

4.6 How a Solutionism mindset may hinder the asking of these questions

“Solutionism” is a term coined by Morozov [66]. One of its definitions is “the belief that all difficulties have benign (usually technological) solutions”. It can be regarded as an outcome of the problem-solving mindset described above, but other issues described under Q1-Q4 also play a role. In their article on lessons learned from successful DSSG projects, Tanweer and

Fiore-Gartland stress the critical importance of expertise on context, project partner organizational culture, and multiple stakeholder perspectives. They conclude that “exposing an inequity or proposing a solution to a social problem doesn’t necessarily mean that social good will follow. If we ignore that warning, we are in danger of lapsing into technological solutionism (Morozov, 2013), where we propose data-informed solutions that have little chance of actually making a difference because they are contextually misconstrued, organizationally untenable, or socially unacceptable.” [17, p. 3]

Table 1 juxtaposes relevant aspects of engineering and social-science mindsets. Solutionism, in this table, is the assumption that a social problem (which usually resides on the right-hand side) is a problem on the left-hand side, coupled with the associated treatment of this problem. The point is not to declare the problem solving approach as useless for AI striving for the Common Good – on the contrary, its clarity and explicitness can often prove highly beneficial for method and system development. Also, AI researchers and practitioners increasingly work in interdisciplinary teams, and sometimes also draw on skilled decision analysts who are much more alert to the complexities of decision making. However, the temptation to consider problems “solved” by a technological “solution” remains strong, and it can stand in the way of seeing and addressing the wider social issues.

5 Is the need to ask these questions not obvious? An exploration of current publications

Some readers may concur that Q1-Q4 are important, but ask: Is it not the case that all computer scientists and AI researchers and practitioners know how important problem formulation is? Is it not the case that they are all aware of the central role of stakeholders, and different stakeholders, in software requirements engineering? Is it not the case that AI researchers and practitioners know about limitations of knowledge, and that computer scientists, coming from a field that has its roots in cybernetics, are aware of systems and their dynamics?

To get a first indication as to how AI in the interest of the Common Good deals with these aspects, I turned to four major venues for AI / DS for (Social) Good. The reasons for this choice are the same as for their choice as providing definitions, explained in Section 2.3. In these conferences, a large number of impressive methods and projects were described. The following analysis in no way intends to downplay these approaches' positive contributions, and caveats with regards to the study's results (which may derive from the conferences' goal being "Social Good" rather than "the Common Good") will be described in the discussion in Section 5.3.

5.1 Surveyed materials

I consulted all hyperlinked contributions (articles, extended abstracts, and presentations) that are made available on the websites of the four venue's most current editions or in their published proceedings. This procedure gave rise to 24 extended abstracts or articles from the Data Science for Social Good Conference 2017²⁸, 4 articles from SoGood 2017 [67], 15 articles from the 2017 AAAI Spring Symposium on AI for Social Good [68], and 56 presentation slide sets, extended abstracts or articles from the 2nd AI for Good Global Summit²⁹. The selection contained all hyperlinked contributions to the second, third and fourth of these venues, since they all presented projects or methods (introductory greetings and other organiza-

tional materials were not considered further). The first venue, the Data Science for Social Good Conference, had multiple tracks, of which three appeared pertinent to the present questions and were therefore analyzed: *DSSG Fellowship Project Talks*, *Short Talks: Research Challenges in Doing Data Science for Social Good* and *Short Talks: Collaboration Models for Social Good*. All contributions were read and assessed with regard to content and whether they contained explicit information relating to the four lead questions above.

5.2 Summary of the findings

5.2.1 Content

The contributions covered a wide range of issues, with no specific issue covered by more than one article. The motivation was usually framed as a problem to be solved; in some cases, a general social-good goal was named or could easily be inferred from the introduction and the specific computational goal. The issues were, in the large majority of cases, of a *substantive* nature. To the best of my knowledge, no generally agreed-upon ontology of Social Good objectives exists; I have therefore performed a rough classification by the SDG goals that were chosen as the guiding principle of the AI for Good Global Summit, and will report only the most frequent ones. I will in some cases distinguish between contributions to the first three venues and those to the fourth, for two reasons: the Summit contributions were mostly slidesets that presented little method detail (such that counts of some methodological questions not being covered may be misleading); and the Summit was structured into four thematic tracks (such that counts of topics in this conference follow from this structure, which is not the case in the other conferences).

Sustainable Development Goal (SDG) #3, Good Health and Well-being, was the most frequent: Eight out of the 43 contributions to the first 3 venues covered topics related to this SDG. The topic was also covered in eleven of the contributions to the fourth venue, in which *AI + Health: Artificial Intelligence – a game changer for Universal Health Coverage?* was one of the four tracks into which the conference was

structured. Five contributions in venues 1-3 could be linked to SDG #11, Sustainable Cities and Communities, and two to transportation (which concerns both cities and SDG #9, Industry, Innovation and Infrastructure). In addition, a second track in the AI for Global Good Summit with 13 contributions focused on *Smart Cities*. The AI for Global Good Summit dedicated a third track, with 6 contributions, on *AI and Satellite Imagery* and linked this method-centric topic specification for contributions to three further SDGs (No Poverty, Life on Land, and Zero Hunger), leading to a strong representation of these topics.

Only a minority of the contributions focused on *procedural* issues. Of these, five were concerned with the scientific process as such, such as crowdsourcing a health-related task to citizen scientists. These contributions already start from the assumption that the goal of the respective scientific project is a legitimate goal for the social/common good (“the aim is not to promote user interaction but to collect useful data for their scientific goals” [69]). The fourth track of the AI for Good Global Summit, *Trust in AI*, in its descriptions and contributions likewise considers the goodness of AI as a given and the need to build trust as a way to convince people of this. Four contributions made proposals for the processes of working towards the Social Good, ranging from DSSG projects via non-profits to UN agencies tasked with identifying SDGs in national development plans.

Three contributions dealt explicitly with democratic processes (in which citizens deliberate about their visions on the Common Good): one presenting a case study platform to create a democratic city planning system [70], one presenting a case study platform to help make city growth equitable by increasing transparency and accountability [71], one proposing an agent-based architecture to predict the effects of policies [72], and one presenting a mathematical voting model [73].

Another goal that could be linked to processes was to improve information access/diffusion and quality (five contributions): the proposed analysis methods for summarizing news and social media contents and identifying misinformation, can arguably help citizens make better-informed democratic choices.

5.2.2 Q1: Alternative goals?

The vast majority of contributions worked with one (sometimes vaguely described or even only implicit) social goal and one computational goal. Thus, there was in general no attempt at framing the social problem in different ways (see Q1 above), and no discussion of whether and how the selected social goal could translate into different computational goals. The difference between the two types of problems (Q1') was not the topic of any paper.

5.2.3 Q2: Different stakeholders, and a description of how their perspectives and needs were assessed?

The large majority of contributions took their problem definition from the non-academic project partner, which was usually also the data provider (for example, a city council, a health agency, or a transportation agency). In some cases, several partners were mentioned (e.g., a tourism agency and a transportation agency), but no conflicts of goals or problems were reported. Tanweer and Fiore-Gartland, in their report on best practices, mentions stakeholders repeatedly, differentiating between “partner organizations” and “affected communities”: “DSSG projects can be more effective when done with consideration for the structures and cultures of partner organizations. This knowledge is often tacit for those stakeholders [...] but without it, DSSG teams run the risk of developing products and services that have little chance of being embraced by stakeholders [...] DSSG teams need to view social issues from multiple perspectives, realizing that different communities and interest groups have [different and] sometimes conflicting stakes in the way social problems are portrayed and addressed. Without understanding the complex political landscapes and contested histories within which social problems are enmeshed, they run the risk of alienating affected communities” [17, p. 3].

Only one contribution described an explicit multi-stakeholder process that was used to formulate the computational/engineering problem [74], and another one mentions that their project partners fol-

lowed such a process, in which there are clearly visible different positions [71]. An agent-based architecture to reflect different positions and interests was proposed in [72], but the question of how to elicit these positions and interests was left implicit in this paper. In two contributions, differences between stakeholders' interests are identified, but only one position is then pursued in the method or tool [75], or the problem is delegated by proposing that the owner of the AI device chooses the position that the machine will follow [76]. Two contributions dealing with questions of fairness take the existence of different viewpoints of what "fair" means as the starting points of their formal models [77, 78]. The contributions dealing with democratic processes, especially [70], implicitly acknowledge different viewpoints, but do not provide any specifics.

5.2.4 Q3: The role of knowledge

It is difficult to describe the breadth and depth of knowledge that was *brought to* the contributions, since that would require an in-depth understanding of all the domains of all the papers, or at least a validated bibliometric method. Both of these are beyond the possibilities of the current article. However, it can be observed that the setup of the venues strongly encourages the participation by AI researchers and practitioners, if only because the venues are defined in a discipline-centric way ("AI for ...", "Data Science for ..."). This limits the incentives for people from other fields to participate.

Concerning the role of knowledge as an *output of* the contribution, the texts give clearer indications. First, data science methods were not only the subject of the "Data Science for ..." venue, but also of many contributions to the "AI for ..." venues (solely or, for example, in combination with computer vision in the AI and Satellite Imagery Track of the AI for Good Global Summit). This leads to a strong representation of knowledge-centric methods. Second, nevertheless, eleven contributions are coupled with an explicitly identified and specific intervention, such as apps designed to detect health problems [79] or apps to incentivize people to cycle to work [80]. A further contribution mentions that the project part-

ner intends to use the developed tool for a number of specified purposes [81]. Two contributions [82, 17] describe collaboration processes and thereby go beyond knowledge. For several contributions in the AI for Good Global Summit, it was difficult to see from the slide-set presentation what roles knowledge and interventions played. Thus, the numbers given here are likely to be a lower bound on those contributions that went beyond knowledge.

5.2.5 Q4: Dynamics

Auerbach et al. [71] interleave data analysis and policy intervention and build a tool to satisfy various information needs in this process. They include a discussion on how insights from their data analysis and possible actions based on these insights could interact in the future and what this would imply for applicants and their use of the financing instrument they study. In the studied sample of contributions, this was the only one that included an explicit consideration of possible dynamics.

5.3 Discussion

The results indicate that so far, the considerations presented in the current paper are not an integral part of current operational research practices in AI/DS for (Social) Good. On the other hand, the importance of Q1 and Q2 is stressed in many methodological papers in the analyzed publications sample. A question akin to Q3 is also highlighted from within the DSSG community: the importance of broader, and experiential, knowledge. Q4 is, at the moment, mostly reflected in the fairness/non-discrimination literature. On the other hand, the mutual dependencies between technology and society are an integral part of the literature on socio-technical systems. More interaction with this field could benefit future AI and data studies [13]. So the community finds a stronger reflection of process relevant, but also finds it hard to translate these into concrete research practices.

Some caveats and encouraging recent developments should be taken into account when interpreting these findings.

Q1 and Q2 were weakly represented. One reason for this may be that the requirements on Q1 and Q2 are likely to be less stringent for *Social Good* than for the *Common Good*. It appears from the definitions put forward by this community (see Section 2.3) as well as the conference survey (see Section 5.2) that Social Good may well be produced by considering only or mainly one, potentially very specific, stakeholder group (such as poultry farmers in Africa [83] or citizens registering as unemployed in one city [84]). Developing AI for these groups and/or relevant use cases involving them may still require researchers to consider various perspectives and problem versions, but the scope is much more limited than the “for all” (members of a given community) of the Common Good. According to some definitions of the Social Good, it also appears legitimate to outsource the (social) problem definition to, for example, an NGO or government agency.

The weak representation of Q1 and Q2 may also be a consequence of research practices. When researchers depend on the collaboration of a project partner (for example, to have access to data or to stakeholders), they may face difficulties if they conceptualize the problem in a way that contradicts the project partner’s notion. This expectation may discourage them from exploring other conceptualizations of the problem.

Regarding the breadth of knowledge brought to the research process (one aspect of Q3), a conference not surveyed here made an interesting decision: the “Fairness, Accountability and Transparency in Machine Learning” (FATML) workshop organizers decided to host, as of 2018, a conference called FAT* and to turn FATML into one sub-event. This decision contributed to a more multidisciplinary perspective on fairness, accountability and transparency than in the years before, with contributions drawing on a more diverse set of stakeholders and problem formulations. As a member of the Steering Committee of FAT*, I am probably biased to see this conference as a success, but the case shows that this widening of scope is possible and can be highly successful in terms of the number and quality of submissions and attendance rates.

The method of the conference survey has limi-

tations. The coding exercise was, by design, exploratory, and the method simple. In future work, a codebook and more coders will be employed. In addition, the results of the coding exercise also suggest that the widespread absence of the considerations Q1-Q4 in the publications may be (partially) an artifact of publication conventions that favor unambiguity and the appearance of a linear and smooth research process. In addition, many of the surveyed documents were very short and therefore concentrated on telling a simple story; a longer paper may have given room to alternatives considered and other details of the research process. As a result, I expect that qualitative interviews with project participants may yield more information.

6 Conclusions: towards ethics pen-testing

This section concludes the article by deriving recommendations for design, and by discussing the scope of the proposals made.

6.1 Recommendations and ethics pen-testing

From the previous considerations, some recommendations can be derived for AI researchers and practitioners who want to contribute to the Common Good:

1. Study the various faces of knowledge and non-knowledge, be wary of framing
2. Identify and involve stakeholders throughout
3. Keep looking for individual and common effects; don’t forget the socio-political
4. Employ proportionality thinking: suitable & necessary, balance of interests?
5. Focus on whole systems instead of parts (e.g. algorithms)
6. Consider feedback loops and causal dynamics

7. Draw on the state of the art: e.g. anonymization, discrimination-aware methods
8. Ask questions, early and again
9. Be ready to be wrong, and embrace learning!

These recommendations follow from the discussions of the four questions and their interlinkage, i.e. in particular from Q3 (1. and 4.), Q1 and Q2 (2.), Q2 (3. and 4.) Q4 (5. and 6.), and the concept of provocations as such (8. and 9.). I added 7. as a general design principle learned from the discussions around the GDPR, which, to protect the rights and freedoms of individuals affected by (data processing) technology, requires technology and organizational designers to design based on the state of the art (see [85] for a discussion of the challenge of translating this legal requirement into engineering practice). This requirement appears pertinent and useful to all design “for Good”.

But how can we ensure that these questions, which are necessarily somewhat vague and “social-science”, actually get asked, and their answers given attention? For this, I propose to draw on an established technique from a core computer-science field, IT security. This method can also be regarded as transposing humanistic self-reflective and self-critical thinking into computer-science mindsets. Specifically, the proposal is inspired by pen-testing in IT security. A penetration test, colloquially known as a pen test, is an authorized simulated attack on a computer system that looks for security weaknesses, potentially gaining access to the system’s features and data. It is well-known that no system is 100% secure. The point is not to pass all pen-tests (this is impossible), the point is to get better through trying. Since IT security researchers and professionals know this, pen-testing is considered valuable and integral to development. Analogous thinking is applied in other subfields of computer science: cryptography and anonymization (where attack modeling is key to any formal proposal), and adversarial machine learning. I propose to apply this mindset in AI design and call the resulting method ethics pen-testing. This consists of asking the lead questions described above with regard

to a proposed design, with the intention of “attacking” the good intentions, the claim to a contribution to the Common Good – as a critical and adversarial method for identifying weaknesses not in order to “fix” the design or make it unequivocally good (because this is impossible), but in order to make it better. Such ethics pen-testing should be carried out not by the researcher/developer themselves, but by others, including representatives of various stakeholder groups. Drawing on techniques from various disciplines, we have sketched the organization of such “attacks” as “tool clinics” [86].

6.2 How specific is this article to AI, and to the goal of the Common Good?

This article was motivated both by the challenges for AI and the specific current concern within AI about the Good and the Common Good. The examples and proposed countermeasures draw specifically on AI, often DS, methods and applications. But is, or in what sense(s) is it specific to AI?

As the discussions of data science and robot(ic)s has shown, it is usually not straightforward to analyze, let alone build a “purely AI” system. Deployed systems, and thus those with real-world ethical implications, usually involve components from many other fields of computer science, such as user interfaces, network functionalities, or databases. The question is then which of these components is “responsible” for real-world effects. For example, are the often-discussed filter bubbles, fake news, and potential manipulation problems a consequence of the AI deployed in social media, or of the way online social networks are connected and information is spread over the Internet? I have heard AI colleagues protective of the supposedly value-free nature of their subfield blame the networking features, and network researchers blame the AI. Conversely, positive effects tend to be attributed to one’s “own” field. Any researcher who uses the lens of socio-technical systems studies would of course observe immediately that indeed these various components together constitute the IT system that creates the effects, and they would

point to the effects of the larger socio-technical systems in which the IT system is embedded.

For these reasons, neither the conceptual analysis presented in this paper, nor the normative statements derived from it, nor the intervention sketched as an approach to operationalizing these statements, can be limited to AI in a clean analytical way. In particular, the problem-solving mindset and the difficulties of stakeholder involvement (lead questions 1 and 2) are integral to computer scientists' and other engineers' professional mindset and practice. To some extent, all interventions share these problems; the term "social engineering" emphasizes this commonality. I suspect that the reductionist tendencies in conceptualizing problems and solutions are stronger among computer-science engineers, but this presumption should be investigated in future work.

Extending the same reasoning, one could apply the "adversarial" method of ethics pen-testing also in a much wider range of settings. For example, methods for constructive technology assessment (CTA) [87] involve multi-perspective criticism and (friendly) "attacks". The question how similar or dissimilar they are to pen-testing should be the subject of future research.

With regard to lead questions 3 and 4, AI is particularly implicated due to its special relationship with knowledge (lead question 3) and on account of the specific side effects and dynamics (lead question 4) of big data, which is a key driver for AI.

In sum, the issues and recommendations of the present paper are characteristic of, but not limited to, artificial intelligence and data science.

Another question of scope concerns the ethical goal: Are the questions and methods proposed here specific to the Common Good? I believe that they are applicable for other goals too, but the questions to be asked will depend on the specific technology fields and design goals. Thus, for example, design in the interest of one specific right or value could also profit from asking the questions and using the recommendation checklist. But the scope may be smaller, and conflicting viewpoints may not need to be considered, or considered only as a background given, in system design when the focus is on one or a small number of stakeholder groups or goals. The interactive system

designed together with an anti-displacement NGO in [71], already referred to in the conference survey, is a good example of a specific focus in stakeholders, goal, and background.

6.3 A road ahead

The idea of ethics pen-testing faces one key challenge: that the testers be perceived as "nagging" and the tested researcher/developer feel attacked in their personal good intentions. Both concerns have been voiced towards me often in earlier presentations. It needs to be kept in mind that in security pen-testing, everybody agrees on what the core goal is and when security is "broken", whereas the reason why one stakeholder considers a system "broken" in ethics pen-testing may be a goal that the designer does not even agree with (which implies that there may be debates on whether the system is broken at all). These challenges can only be overcome if researchers and practitioners take the ethical quality of their products as seriously as the formal and engineering qualities, and if they regard this as a professional rather than personal virtue.

Finally, let us not forget that we need to attack the choice of reflective questions and the idea of ethics pen-testing itself too – theoretically and on the basis of case studies. This is an invitation to the AI community to test and share their experiences!

Acknowledgements

This article arose from a keynote given at the March 2017 AI in Asia – AI for Social Good workshop in Tokyo, Japan, and the great discussions we had there. I am also indebted to the participants of the December 2017 workshop on Intersectionality and Algorithmic Discrimination at the Lorentz Centre, Leiden, The Netherlands. I thank Doris Allhutter, Robert Kahlert, Geoffrey Rockwell, Salvatore Ruggieri and Roger Vergauwen for valuable comments on earlier drafts.

References

- [1] L. Pangrazio, Exploring provocation as a research method in the social sciences. *International Journal of Social Research Methodology*, 20 (2) (2017), 225
- [2] d. boyd, K. Crawford, Critical questions for big data. *Information, Communication & Society* 15 (5) (2012), 662
- [3] ACM (1992). ACM Code of Ethics and Professional Conduct. <https://www.acm.org/about-acm/acm-code-of-ethics-and-professional-conduct>
- [4] Future of Life Institute. Asilomar Principles (2016). <https://futureoflife.org/ai-principles/>
- [5] IEEE (2016). Ethically Aligned Design. A Vision for Prioritizing Human Wellbeing with Artificial Intelligence and Autonomous Systems. Version 1 for public discussion. http://standards.ieee.org/develop/indconn/ec/ead_v1.pdf
- [6] S. Lee, Common Good. In: *Encyclopedia Britannica* (n.d.). <https://www.britannica.com/topic/common-good>
- [7] W. Hussain, The Common Good. In: E.N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*, Spring 2018 edition. <https://plato.stanford.edu/archives/spr2018/entries/common-good/>
- [8] C. Blum, Determining the Common Good: A (re-)constructive critique of the proceduralist paradigm. *Phenomenology and Mind*, 3 (2012), 176
- [9] K. Popper, *The Open Society and Its Enemies* (Routledge, UK, 1945)
- [10] M. Jaede, The Concept of the Common good. Working Paper Series of the Political Settlements Research Programme (PSRP) of the University of Edinburgh (Edinburgh, UK, 2017). <https://www.britac.ac.uk/sites/default/files/Jaede.pdf>
- [11] J. Cohen, Procedure and substance in deliberative democracy. In: J. Bohman & W. Rehg (Eds.), *Deliberative Democracy : Essays on Reason and Politics* (MIT Press, Boston, MA 1997) 407
- [12] G. Capoccia, Militant democracy: The institutional bases of democratic self-preservation. *Annual Review of Law and Social Science*, 9 (1) (2013), 207
- [13] R. De Wolf, E. Vanderhoven, B. Berendt, J. Pierson, T. Schellens, Self-reflection on privacy research in social networking sites. *Behaviour & Information Technology* 36 (5) (2017) 459
- [14] World Economic Forum, Artificial Intelligence for the Common Good. Sustainable, Inclusive and Trustworthy (2018) <https://weforum.ent.box.com/v/AI4Good>
- [15] North Highland Consulting, AI for the Common Good. An Ethical Framework to Harness AI's Greatest Potential (2018) <http://www.northhighland.com/insights/white-papers/ai-for-the-common-good>
- [16] B. Mols, AI For the Common Good. *ACM News*, Jun. 20, 2017 <https://cacm.acm.org/news/218696-ai-for-the-common-good/fulltext>
- [17] A. Tanweer, B. Fiore-Gartland, Cross-sector collaboration in Data Science for Social Good: Opportunities, challenges, and open questions raised by working with academic researchers. In: *Data Science for Social Good Conference 2017*, Sep. 28-29, 2017, Chicago, IL <http://dssg.uchicago.edu/wp-content/uploads/2017/09/tanweer.pdf>
- [18] B. J. Copeland, Artificial intelligence, In: *Encyclopedia Britannica* (n.d.) <https://www.britannica.com/technology/artificial-intelligence>
- [19] L. Cao, Data Science: A comprehensive overview. *ACM Computing Surveys* 50 (3) (2017) Article No. 43

- [20] D. Conway, The Data Science Venn Diagram (n.d.) <http://drewconway.com/zia/2013/3/26/the-data-science-venn-diagram>
- [21] Wikibooks, The Free Textbook Project, Cognitive Psychology and Cognitive Neuroscience/Knowledge Representation and Hemispheric Specialisation (2017) https://en.wikibooks.org/w/index.php?title=Cognitive_Psychology_and_Cognitive_Neuroscience/Knowledge_Representation_and_Hemispheric_Specialisation&oldid=3277633
- [22] S. J. Russell, P. Norvig, Artificial Intelligence: A Modern Approach (Prentice Hall, Englewood Cliffs, NJ, 1995)
- [23] U.M. Fayyad, G. Piatetsky-Shapiro, P. Smyth, Knowledge discovery and data mining: Towards a unifying framework. In: Proc. Int. Conf. on Knowledge Discovery and Data Mining (SIGKDD) (1996) 82
- [24] G. Veruggio, Roboethics [TC Spotlight]. IEEE Robotics & Automation Magazine, 17 (2) (2010), 105
- [25] P. M. Asaro, What should we want from a robot ethic? International Review of Information Ethics, 6 (2006), 9
- [26] B. F. Malle, Integrating robot ethics and machine morality: The study and design of moral competence in robots. Ethics and Information Technology, 18 (2016), 243
- [27] W. Wallach, C. Allen, Moral machines: Teaching robots right from wrong (Oxford University Press. New York, NY, 2008)
- [28] S. A. Fricker, R. Grau, A. Zwingli, Requirements engineering: Best practice. In: S.A., Fricker, C. Thummler, & A. Gavras (Eds.), Requirements Engineering for Digital Health (Springer, USA 2015) 25
- [29] AOIR (Association of Internet Researchers), Ethical Decision-Making and Internet Research: Recommendations from the AoIR Ethics Working Committee (Version 2.0). (2012) <https://aoir.org/reports/ethics2.pdf>
- [30] J. Goguen, Requirements engineering as the reconciliation of technical and social issues. In: Requirements Engineering: Social and Technical Issues (Academic Press 1994) 165
- [31] S. M. Bird, A. McAuley, S. Perry, C. Hunter, Effectiveness of Scotland's National Naloxone Programme for reducing opioid-related deaths: a before (2006–10) versus after (2011–13) comparison. Addiction, 111 (2016) 883
- [32] I. Amrani, Here in Colombia, the hypocrisy of western cocaine users is laid bare. The Guardian, Aug. 1, 2018, <https://www.theguardian.com/commentisfree/2018/aug/01/colombia-hypocrisy-western-middle-class-cocaine-users-violence>
- [33] A. Kovacevic, Engineering Design Process – Part 1. Problem Definition (2017) <http://www.staff.city.ac.uk/~ra600/ME1105/Lectures/ME1110-11.pdf>
- [34] P. Naur, Formalization in program development. BIT, 22, 437
- [35] F. Villamor, Death of Philippine teenager stokes opposition to Duterte's drug crackdown. The New York Times, Aug. 23, 2017, <https://www.nytimes.com/2017/08/23/world/asia/duterte-drug-crackdown.html?mcubz=3>
- [36] ACLU, Written Submission of the American Civil Liberties Union on Racial Disparities in Sentencing. Hearing on Reports of Racism in the Justice System of the United States submitted to the Inter-American Commission on Human Rights (2014) https://www.aclu.org/sites/default/files/assets/141027_iachr_racial_disparities_aclu_submission_0.pdf
- [37] K. Gwynne, 4 biggest myths about crack. Salon, Oct. 8, 2013, https://www.salon.com/2013/10/08/4_biggest_myths_about_crack/

- [//www.salon.com/2013/08/10/busting_the_crack_propaganda_myths_partner/](http://www.salon.com/2013/08/10/busting_the_crack_propaganda_myths_partner/)
- [38] J. Mechanic, When a drug epidemic hit white America, addiction became a disease. *Huffington Post*, Jul. 10, 2017, https://www.huffingtonpost.com/entry/when-a-drug-epidemic-hit-white-america-addiction-became_us_5963a588e4b08f5c97d06b9a
 - [39] K. Lum, W. Isaac, To predict and serve? Significance, 13(5), 14
 - [40] A. Chouldechova, Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big Data*, 5(2) (2017), 153
 - [41] J. Kleinberg S. Mullainathan, M. Raghavan, Inherent trade-offs in the fair determination of risk scores. In: *Proc. 8th Conf. on Innovations in Theoretical Computer Science (ITCS)*, 2017
 - [42] D. Tewksbury, D.A. Scheufele, News framing theory and research. In J. Bryant, M. B. Oliver (Eds.), *Media effects: Advances in theory and research* (Earlbaum, Hillsdale, NJ, 2009) 17
 - [43] D. Baum, Legalize it all. How to win the war on drugs. *Harper's Magazine*, Apr. 4, 2016, <https://harpers.org/archive/2016/04/legalize-it-all/>
 - [44] Wikipedia contributors, David Nutt, In: *Wikipedia, The Free Encyclopedia* (2018) https://en.wikipedia.org/w/index.php?title=David_Nutt&oldid=852024763
 - [45] S. Ewen, *PR! A Social History of Spin.* (Basic Books, New York, 1996)
 - [46] E. Bakshy, S. Messing, L.A. Adamic, Exposure to ideologically diverse news and opinion on Facebook. *Science*, 348 (2015), 1130
 - [47] M. Del Vicario, G. Vivaldo, A. Bessi, F. Zollo, A. Scala, G. Caldarelli, W. Quattrociocchi, Echo chambers: Emotional contagion and group polarization on Facebook. *Scientific Reports*, 6 (2016), Article 37825
 - [48] H. T. P. Williams, J. R. McMurray, T. Kurz, F. H. Lambert, Network analysis reveals open forums and echo chambers in social media discussions of climate change. *Global Environmental Change*, 32 (2015), 126
 - [49] T. T. Nguyen, P.-M. Hui, F. M. Harper, L. Terveen, J. A. Konstan, J.A., Exploring the filter bubble: The effect of using recommender systems on content diversity. In: *Proc. WWW 2014* (ACM, New York, 2014).
 - [50] L. Taylor, The ethics of big data as a public good: which public? Whose good? *Phil. Trans. R. Soc. A*, 374 (2016), 20160126.
 - [51] L. Taylor, Safety in numbers? Group privacy and big data analytics in the developing world. In: L. Taylor, L. Floridi, B. van der Sloot (Eds.), *Group Privacy, Philosophical Studies Series 126* (Springer, Berlin etc., 2017) 13
 - [52] C. Barabas, K. Dinakar, J. Ito, M. Virza, J. Zittrain, Interventions over predictions: Reframing the ethical debate for actuarial risk assessment. *Proceedings of Machine Learning Research*, 81 (2018) <http://proceedings.mlr.press/v81/barabas18a/barabas18a.pdf>
 - [53] M. Kunaver, T. Požrl, Diversity in recommender systems – a survey. *Knowledge-Based Systems*, 123 (2017), 154
 - [54] B. Berendt, B. Gao, S. Gürses, T. Peetz, J. Pierson, *SPION Deliverable 5.2 – Report on Research Activities (Feedback and Awareness Tools).* COSIC Technical Report, KU Leuven, (Leuven, Belgium, 2014) <https://www.cosic.esat.kuleuven.be/publications/article-2496.pdf>
 - [55] A. Jameson, B. Berendt, S. Gabrielli, F. Cena, C. Gena, F. Vernerio, K. Reinecke, Choice architecture for human-computer interaction. *Foundations and Trends in Human-Computer Interaction*, 7 (1-2) (2014), 1
 - [56] R. Shamir, The age of responsabilization: on market-embedded morality. *Economy and Society*, 37 (1) (2008), 1

- [57] G. Rockwell, B. Berendt, Information wants to be free, or does it? The ethics of datafication. *Electronic Book Review* (2017) <http://electronicbookreview.com/thread/technocapitalism/datafiction>
- [58] M. Gross, L. McGoey (Eds.) *Routledge International Handbook of Ignorance Studies*. (Routledge, London / New York, 2015)
- [59] A. Yakushev, S. Mityagin, Social networks mining for analysis and modeling drugs usage. In *Proc. Of ICCS 2014* (2014)
- [60] M. Kosinski, D. Stillwell, T. Graepel, Private traits and attributes are predictable from digital records of human behavior. *PNAS*, 110 (15) (2013), 5802
- [61] P. Greenfield, The Cambridge Analytica files: The story so far. *The Guardian*, Mar. 26, 2018, <https://www.theguardian.com/news/2018/mar/26/the-cambridge-analytica-files-the-story-so-far>
- [62] S. Barocas, A. D. Selbst, Big data's disparate impact. *104 California Law Review*, 671 (2016). <http://dx.doi.org/10.2139/ssrn.2477899>
- [63] C. O'Neil, *Weapons of Math Destruction*, (Crown Publishers, New York, 2016)
- [64] B. Berendt, S. Preibusch, Toward accountable discrimination-aware data mining: The importance of keeping the human in the loop – and under the looking-glass. *Big Data*, 5 (2) (2017), 135
- [65] D. Ensign, S. A. Friedler, S. Neville, C. Scheidegger, C., S. Venkatasubramanian, Run-away feedback loops in predictive policing. *Proceedings of Machine Learning Research*, 81 (2018) <http://proceedings.mlr.press/v81/ensign18a/ensign18a.pdf>
- [66] E. Morozov, *To Save Everything, Click Here: The Folly of Technological Solutionism*. (Public Affairs, New York, 2013)
- [67] R. Gavaldà, I. Koprinska, S. Kramer (Eds.), *Proceedings of the Second Workshop on Data Science for Social Good co-located with European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD 2017)*, Skopje, Macedonia, Sep. 18, 2017, *CEUR Workshop Proceedings Vol-1960*, <http://ceur-ws.org/Vol-1960/>
- [68] E. Horvitz et al. (Eds.), *Artificial Intelligence for the Social Good. Papers from the 2017 AAAI Spring Symposium*. Technical Report SS-17-01, (AAAI, Palo Alto, CA, 2017) <https://www.aaai.org/Library/Symposia/Spring/ss17-01.php>
- [69] J. Garriga, J. Piera, F. Bartumeus, A Bayesian framework for reputation in citizen science. In: [67], <http://ceur-ws.org/Vol-1960/paper6.pdf>
- [70] F. Bria, The role of cities in democratizing AI and data ownership: Learning from Barcelona. Presentation at The AI for Good Global Summit 2018, Geneva, May 15-17, 2018, <https://www.itu.int/en/ITU-T/AI/2018/Documents/Presentations/Francesca%20Bria.pdf>
- [71] J. Auerbach, H. Barton, T. Blunt, V. Chaganti, B. Ghai, A. Meng, C. Blackburn, E. Zegura, P. Flores, Coupling data science with participatory planning for equity in urban renewal programs: An analysis of Atlanta's Anti-Displacement Tax Fund. In: *Data Science for Social Good Conference 2017*, Sep. 28-29, 2017, Chicago, IL <http://dssg.uchicago.edu/wp-content/uploads/2017/09/auerbach.pdf>
- [72] V. Dignum, F. Dignum, Societal challenges need social agents. In: [68], <https://aaai.org/ocs/index.php/SSS/SSS17/paper/view/15302>
- [73] M. Prasad, Back to the future: A framework for modelling altruistic intelligence explosions. In: [68], <https://aaai.org/ocs/index.php/SSS/SSS17/paper/view/15326>

- [74] Y.-H. Wang, Y.-Y. Chen, S.-C. Chen, C.-K. Liu, T.C. Hsieh, Data for Social Good: A case study of building an effective public-private partnership on domestic violence prevention. In: Data Science for Social Good Conference 2017, Sep. 28-29, 2017, Chicago, IL <http://dssg.uchicago.edu/wp-content/uploads/2017/09/hsieh.pdf>
- [75] A. Delaunay, J. Guérin, Wandering detection within an embedded system for Alzheimer suffering patients. In: [68], <https://aaai.org/ocs/index.php/SSS/SSS17/paper/view/15317>
- [76] O. Bendel, LADYBIRD: The animal-friendly robot vacuum cleaner. In: [68], <https://aaai.org/ocs/index.php/SSS/SSS17/paper/view/15277>
- [77] P. Baumgartner, Challenges in assessing predictive bias. In: Data Science for Social Good Conference 2017, Sep. 28-29, 2017, Chicago, IL <http://dssg.uchicago.edu/wp-content/uploads/2017/09/baumgartner.pdf>
- [78] K. Gummedi, A. Weller, Cross-cultural perceptions of fairness in algorithmic decision making: A case study of criminal risk prediction. Presentation at The AI for Good Global Summit 2018, Geneva, May 15-17, 2018, <https://www.itu.int/en/ITU-T/AI/2018/Documents/Presentations/Gummedi%20and%20Weller.pdf>
- [79] I. Weber, Using hyper-targeted advertising for public health messaging. Presentation at The AI for Good Global Summit 2018, Geneva, May 15-17, 2018, <https://www.itu.int/en/ITU-T/AI/2018/Documents/Presentations/Ingmar%20Weber.pdf>
- [80] E. Santo, UMO. Smarter Cities. Happier people. AI for *real* urban sustainability. Presentation at The AI for Good Global Summit 2018, Geneva, May 15-17, 2018, <https://www.itu.int/en/ITU-T/AI/2018/Documents/Presentations/Eyal%20Santo.pdf>
- [81] M. Chappelka, J. Oh, D. Scott, M. Walker-Holmes, Food for thought: Analyzing public opinion on the supplemental nutrition assistance program. In: Data Science for Social Good Conference 2017, Sep. 28-29, 2017, Chicago, IL <http://dssg.uchicago.edu/wp-content/uploads/2017/09/scott.compressed.pdf>
- [82] C.-K. Liu, T. C. Hsieh, Lessons learned from using data science to empower change agents across data silos. In: Data Science for Social Good Conference 2017, Sep. 28-29, 2017, Chicago, IL <http://dssg.uchicago.edu/wp-content/uploads/2017/09/liu.pdf>
- [83] E. Nwankwo, Building trust with East African farmers: A poultry app for Good. Presentation at The AI for Good Global Summit 2018, Geneva, May 15-17, 2018, <https://www.itu.int/en/ITU-T/AI/2018/Documents/Presentations/Dina%20and%20Ezinne.pdf>
- [84] H. Holmestad, Predicting risk of long-term unemployment. In: Data Science for Social Good Conference 2017, Sep. 28-29, 2017, Chicago, IL <https://dssg.uchicago.edu/wp-content/uploads/2017/09/holmestad.pdf>
- [85] S. Schiffner et al., Towards a roadmap for privacy technologies and the General Data Protection Regulation: A transatlantic initiative. In: Proceedings of the Annual Privacy Forum 2018, Jun. 13-14, 2018, Barcelona (Springer, Berlin etc., 2018)
- [86] A. Morton, B. Berendt, S. Gürses, J. Pierson, “Tool Clinics”: Embracing multiple perspectives in privacy research and privacy-sensitive design. Dagstuhl Reports, 3 (7) (2013), 96
- [87] A. Genus, Rethinking constructive technology assessment as democratic, reflective, discourse. Technological Forecasting and Social Change, 73 (1) (2006), 13
- [88] H. Hanson, Nixon aides suggest colleague was kidding about drug war being designed

to target black people. The Huffington Post, Mar. 25, 2016, http://www.huffingtonpost.com/entry/richard-nixon-drug-war-john-ehrllichman_us_56f58be6e4b0a3721819ec61

- [89] Wikipedia contributors, Nuremberg Code, In: Wikipedia, The Free Encyclopedia (2018) https://en.wikipedia.org/w/index.php?title=Nuremberg_Code&oldid=848155807
- [90] N. A. Patel, G. D. Elkin, Professionalism and conflicting interests: The American Psychological Association's involvement in torture. *AMA Journal of Ethics*, 17 (10) (2015), 924

Notes

¹<https://en.oxforddictionaries.com/definition/ethical>. The term is often used, also in AI publications, in this everyday meaning rather than in the scientific meaning of relating to ethics.

²All URLs referenced in this article were last retrieved on August 16, 2018.

³The term is also used in different disciplines. For example, in economics, something is a common good if no-one can be excluded from “consuming the good”, regardless of whether it is a benefit or not. This concept relates to issues of access, distribution and scarcity of resources. The discussion of these issues would go beyond the scope of this article.

⁴<https://dssg.uchicago.edu/>. This description has been used since the first version of the site recorded in the Web Archive at <https://web.archive.org/web/20140318063609/http://www.dssg.uchicago.edu:80/>

⁵from the first edition at <https://sites.google.com/site/ecmlpkdd2016sogood/>

⁶<https://www.aaai.org/Library/Symposia/Spring/ss17-01.php>

⁷<https://aaai.org/Symposia/Spring/sss17symposia.php>

⁸<http://scf.usc.edu/~amulyaya/AISOC17/>

⁹<https://www.itu.int/go/AIforGood2018>, <https://sustainabledevelopment.un.org/?menu=1300>

¹⁰<https://www.fatml.org/>, <http://fatconference.org/>

¹¹<https://www.acm.org/publications/class-2012>

¹²Cultural diversity was not mentioned explicitly in the 1948 UDHR, but is implicit in its Article 27 when interpreted together with the 2001 UNESCO Declaration on Cultural Diversity, see http://portal.unesco.org/en/ev.php-URL_ID=13179&URL_DO=DO_TOPIC&URL_SECTION=201.html.

¹³Richard Nixon's statement about drug abuse in a famous press conference 1971, see transcript at <http://www.presidency.ucsb.edu/ws/?pid=3047>

¹⁴Donald Trump's declaration of 2017, repeated in 2018, see <https://www.whitehouse.gov/briefings-statements/president-donald-j-trump-combatting-opioid-crisis/>.

¹⁵see transcript at <https://www.whitehouse.gov/briefings-statements/remarks-president-trump-combatting-opioid-crisis/>

¹⁶Of course, it is a simplification to equate such framings with actors; different interests, values, norms and positionings of the problem were included, because others are outside the scope of an AI approach. Number 7 illustrates the contextual nature of problem definition: This may be an addict's view when on withdrawal, even if the same person would under other circumstance worry more about the negative effects on health and life prospects.

¹⁷“3.4 Ensure that users and those who will be affected by a system have their needs clearly articulated during the assessment and design of requirements; later the system must be validated to meet requirements.”

¹⁸cf. for example the sections on “Key guiding principles” and on “Internet specific ethical questions”

¹⁹Section 2, Embedding Values Into Autonomous Intelligent Systems

²⁰All definitions of “problem” from <https://en.oxforddictionaries.com/definition/problem>

²¹The distinction between data, information, and knowledge is a classical discussion in AI, but since the focus here is on the practical consequences of deployed AI technology rather than on its philosophical foundations, the discussion is not relevant for present purposes.

²²This has not gone undisputed, see [88].

²³Examples are described at <https://www.addiction.com/12575/six-sobriety-apps-you-should-know-about/>, <https://www.technologyreview.com/s/604085/treating-addiction-with-an-app/>, and <https://saferlockrx.com/top-apps-that-help-parents-prevent-teen-drug-abuse/>

²⁴cited after [89]

²⁵“We must use critical thought to distinguish what is ethical from what is lawful and to consider what it means to be a professional. Therefore, we must continually question and re-question authority, whether it is the law or a code of ethics, or else we may be doomed to serve the interests of those who crafted the code, not necessarily the interests of those who need to embody the code or use it to guide their practice. Just because a principle is codified does not make it ethical.” [90]

²⁶“Multiple judgments are possible, and ambiguity and uncertainty are part of the process. We advocate guidelines rather than a code of practice so that ethical research can remain flexible, be responsive to diverse contexts, and be adaptable to continually changing technologies.” [29]

²⁷Social network platforms and apps, and quantified-self apps (see above), are often not very secure. Data and system security, while key in the GDPR and also mentioned in some AI ethics codes but not in others, does not receive as much attention in AI as it probably should.

²⁸<https://dssg.uchicago.edu/data-science-for-social-good-conference-2017/agenda/>

²⁹<https://www.itu.int/en/ITU-T/AI/2018/Pages/programme.aspx>