Towards an Ethical Code for Data-Based Business

Loi, Michele
Digital Society Initiative
University of Zurich
Zurich, Switzerland
michele.loi@uzh.ch

Heitz, Christoph
Zurich University of Applied Sciences
and Digital Society Initiative, UZH
Winterthur, Switzerland
christoph.heitz@zhaw.ch

Ferrario, Andrea
Mobiliar Lab for Analytics at ETH
ETH Zurich
Zurich, Switzerland
aferrario@ethz.ch

Schmid, Anita
Data Strategy & Science
Migros-Genossenschafts-Bund
Zurich, Switzerland
anita.schmid@mgb.ch

Christen, Markus
Digital Society Initiative
University of Zurich
Zurich, Switzerland
christen@ethik.uzh.ch

***Abstract*—In this paper, we outline the structure and content of a code of ethics for companies engaged in data-based business, i.e. companies whose value propositions strongly depends on using data. The code provides an ethical reference for all people in the organization who are responsible for activities around data. It is primarily targeting private industry, but public organizations and administrations may also use it. A joint industry-academic initiative, involving specialists for ethics as well as for all relevant data-related issues, developed this code.**

Keywords—ethics, data, algorithm, predictive models, fairness, privacy, algorithmic bias, code.

I. Introduction

This paper provides a summary description of a new ethical Code that was recently developed within the Swiss Alliance for Data-Intensive Services [1]. This paper is *not* the code, although it cites some of its contents. A preliminary version of the *Code* is accessible online [2]. The goal of the *Code* is to give practical orientation to create data-based services that meet reasonable ethical standards and, for that reason, are expected to be more aligned with the expectations of clients, employees and society. The structure of the paper is as follows. Section II provides the social background from which the need for a code originates and the possible contribution of such code to data governance. Section III describes the 'state of art' relative to ethical guidelines by emphasizing differences between the *Code* and many others that have been formulated in recent times. Section IV is the methodology section, which describes how the *Code* was built. Section V outlines the structure of the *Code* and presents the key ethical ideas. Section VI discusses the scope, normative premises, and business relevance of the proposed *Code.*

II. Background

As a part of the rapid development of information technology and digitization, more and more data-based services and products are being developed and deployed. This development changes considerably how we live, work, and interact. Indeed, digitization has the potential to transform our society as a whole. Social media platforms like Facebook, for example, have changed the way people connect, how lives are shared among friends and in communities, as well as how and which information and values are received, processed and re-distributed. This has a deep impact on the individual and the social levels that we only gradually begin to understand.

Private companies and their digital services are main drivers of these changes. They change the lives of customers and citizens, corporations, and societies – to the better or to the worse. The Facebook/Cambridge Analytica case [3] that became public in 2016 exemplifies the heated debates regarding the societal impact of such services. Thus, with the decision to offer a data-based service or product, there comes an ethical responsibility: companies should think about the consequences of the use of their new data-based products and services. Typically, the development of a new data-based product or service is fuelled by the hope that its proposition is positively received by customers and creates value for them. However, negative consequences may also arise from their use; and companies have to care about them.

During the last years, the call for a more responsible use of data, or "data ethics", has intensified. There is increasing pressure on companies to adhere to ethical expectations of customers, employees and the society. Self-regulation, including ethical codes, can achieve this, at least in part. However, the *Code* presented in this paper does not exclude and is not meant to pre-empt legal regulation. It includes suggestions that are in line with current legislation. For example, the recommendation "do not store personal data that is not necessary for providing, improving or expanding your service" reflects the data protection principle of data minimization. Like other ethics expert groups, e.g. [4, pp. 2–3], we fully acknowledge that digital systems do not operate in a lawless world. For example, the European General Data Protection Regulation of 2018 [5] is an example of an attempt to ensure that data use conforms with societal values. Similar legislation exists or is currently being pushed forward worldwide. We recommend that ethical businesses should first of all be lawful and if tensions between ethics and the law emerge, the responsibility to align them falls on several different stakeholders, including those involved in the political process. It is, thus, not the goal of these guidelines to substitute current regulation.

However, the aim of an ethical codex is to exceed legal requirements, for example by providing suggestions for ethical data management beyond legal necessity and taking care of those societal tensions around ethics of data-driven products and services that existing norms still fail to intercept. Such role for ethics has been long acknowledged in the field of business ethics [6]. In particular, Carrol's [7], [8] views have been widely influential in supporting the "total social responsibility of business which includes economic, legal, ethical and discretionary categories" [9, p. 60]. In Carrol's view, a firm, just like a physical person, has responsibilities in different dimensions (see Fig. 1). These responsibilities have different levels of stringency: a firm is strictly *required* to be profitable and to obey the law. It is *socially expected* to behave ethically. Philanthropy is *desired* but not expected. Overall, it is undeniable that companies realize more and more that they have to care about the ethical dimensions of their business, if only for reputational reasons.

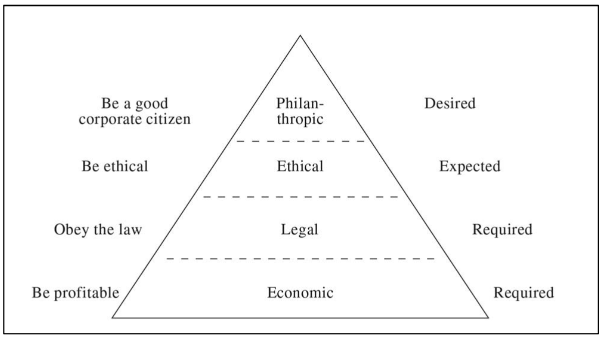

Fig. 1. Carroll's (1991) Pyramid of Corporate Social Responsibility [8]

Only few practical guidelines are available that support data scientists, engineers and product developers to ensure that their products meet the ethical expectations of their stakeholders. Some data ethics codes remain on a rather abstract level and describe general ethical orientations such as values and principles. To the best of our knowledge, no other ethical code organizes ethical prescriptions along the steps of the data pipeline (see Section V) underlying data-based products and services as this one does.

III. State of the Art

Several codes of ethics have emerged from the current multi-stakeholder debate, concerning both data and so-called Artificial Intelligence (AI) [10]–[19]. The final draft of "Ethics guidelines for trustworthy AI," by the High-Level Expert Group on AI of the EU Commission (henceforth EU Guidelines) has been made available while this paper was still under review [4]. Overall, existing codes mention important human values, drawing from human rights, ethical values (well-being, control and autonomy, privacy and intimacy, freedom, shared benefit/prosperity, fairness, equity, equality), procedural values (transparency, explainability, accountability, auditability, responsibility, non-discrimination), and political values (solidarity, democratic participation, sustainability, freedom, common good, justice). Most prominently, the values of human agency, transparency, privacy, fairness and accountability mentioned in the EU Guidelines play an important role informing the contents of these guidelines. This *Code* is sensitive to these values (even though it does not refer to all of them explicitly). Among the existing guidelines that address application issues explicitly, some (e.g. the FAT ML principles [10]) only concern one step of the data process (the design and launch of algorithms). Others (e.g. the Data Ethics Canvas [11]) provide a checklist on *actions*, rather than principles, but do not map these actions into steps in the creation and commercialization of the data-based services or products. The EU Guidelines contain a very concrete "Assessment list" that is meant to guide the application of its guidelines, with checklist items mapping into do's and don'ts of this code. However, to the best of our knowledge, none of the existing codes distinguishes between the different responsibilities associated to different steps in the data pipeline. This makes the implementation of high-level principles and values difficult in real world scenarios. This is why we chose a completely different structure for our *Code*. In fact, instead of structuring the chapters of the *Code* in a top-down way (from values/principles to practical indications and requirements), we start from the practice, e.g. the practice of managing data along a well-defined data pipeline, which underlies most data-based products and services, and unravel the implications of ethical values and principles in that context. The result is a list of *do's and don'ts*, linked to more abstract *key ethical ideas*. These do not represent general ethical principles or values in the sense of philosophical ethics but rather ethical goals, desiderata, requirements, and other value-laden claims *relevant to orient action ethically in a specific context*, or they may be *general facts that would be unethical to disregard*. Thus, our *Code* begins with practice, and aligns all its ethical content to the different steps of working with data. The hoped-for result is that these ethical recommendations and requirements are more easily mapped to different processes, roles and their governance within companies.

IV. Methodology of Code Creation

Members of the Data Ethics Expert Group of the Swiss Alliance for Data-Intensive Services as well as external experts have developed the content of the *Code* in an open deliberation process. In total, 34 persons participated in the process (five different workshops and several conference calls between different subsets of the Expert Group) whose expertise background included ethics, law and technology (data science, ICT, information security and others). Even through the participants are affiliated with particular stakeholder organizations (either academia or Swiss companies), all voices in the code development process were treated as equals.

The development of the code included three steps. In a first step, the examples of existing codes were evaluated and the basic structure of the code was decided (see Section V). In a second step, an editorial board consisting of six people was set up which was responsible for collecting ethical questions, structuring them, proposing concrete guidelines and organizing five review cycles over a period of 11 months within the complete group. This step has led to the current version discussed in this paper. In a third *future* step, a public deliberation process intends to collect feedback both from members of the Data & Service Alliance as well as external stakeholders. Based on this feedback, a refined version of the *Code* will be published by the end of 2019. The full list of contributors will be made available and regularly updated online together with each update of the code text, through the link provided.

V. Structure of the Code

When a large company develops and deploys a data-based service, there are often different roles within the company that are involved, and they have distinct ethical responsibilities. For example, the person who acquires data as a basis for the data-based service is responsible for an ethical behavior in data collection. This includes getting proper consent of the data provider and keeping track of any usage restrictions. The person responsible for managing the data has to make sure that data is stored securely and access is properly restricted. Data scientists who are creating knowledge out of data are responsible for making accurate use of analytical techniques and for drawing correct conclusions from any analytics. Finally, the product manager is responsible for ensuring that the usage of the service does not violate the privacy of its customers, or discriminates between them in an ethically unjustifiable manner. Typically, the decisions with respect to data use are distributed over an organization, and so is the ethical responsibility. For example, a corporate IT department may be responsible for storing data and its management (e.g. data access control) and thus for all ethical aspects of these activities, while product managers are responsible to develop and deploy data-based services and products in the marketplace and thus have to consider the ethical aspects of providing such services. Note that products and services may also be assembled by combining inputs from specialized companies, each of which dealing with one, or a few distinct, tasks of the data-pipeline, thus involving actors of different companies. The challenge of ethical data governance is to identify all data-related ethical issues and to take care of them, sequentially, along the service or product creation pipeline, involving different actors within a company and/or different agents in a data ecosystem. Furthermore, in order to *implement any guidance provided by an ethical code,* it requires a clear definition of roles, responsibilities and monitoring procedures. The latter requirement is especially important because rarely one person alone is able to assess all aspects highlighted above. For example, a data scientist is able to assess the ethical questions around the use of predictive algorithms, but this does not necessarily mean that he can also assess the level of privacy protection for data involved in machine learning. Therefore, for ethical data governance to be successful, data-related decisions involved in the design and deployment of products and services, belonging to different steps in the data pipeline, should be mapped against a common set of ethical and corporate values.

There might be many ways of differentiating between the different tasks or steps in the data pipeline. In fact, multiple process models are used in practice (e.g. CRISP-DM or SEMMA) [20], [21]. For our *Code*, we found that a sequence of four steps is a good choice, as it reflects the typical steps of data management during development and deployment (see Fig. 2) of data-based products and services.

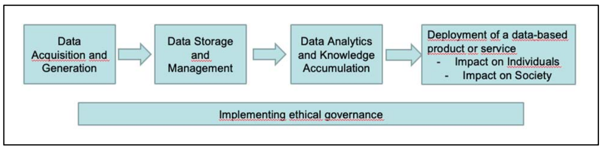

Fig. 2. Main steps in the data pipeline underlying data-based products and services.

For the sake of simplicity, iterations between steps in the data pipeline are not considered. In addition to the four steps of the data pipeline, step 5 is an organizational step, concerned with the creation of ethical data governance. This step obviously relates to all of the preceding, as it relates to the proper organization of the activities implied by the dos and don'ts in each step. We shall refer to this model as '4+1 step' model.

For each of 4+1 steps, we identified two or three key ethical ideas to be considered and we introduced a list of "dos and don'ts". These are concrete and actionable recommendations. The dos and don'ts are comparable, for example, to preliminary checklists items presented the last section of the EU Guidelines. For example. EU Guideline items pertaining to "Privacy and Data Governance" mostly map into dos and don'ts of *Data Acquisition and Generation* and *Data Storage and Management* in the *Code*. Many items listed in the EU Guideline sub-sections concerning "Accuracy", "Reliability and Reproducibility", "Traceability" "Explainability", "Unfair bias avoidance", map into the dos and don'ts of step 3 of the data pipeline, i.e. *Data Analytics and Knowledge Accumulation.* Many EU Guideline items in sub-section "Fundamental Rights" map into our items in step 4-A of the data pipeline (which refers to the *Impact on Individuals*), and items in subsection "Society and democracy" map into our items of step 4-B (*Impact on Society*). Many EU Guideline items in sub-sections *Traceability, Minimizing and Reporting Negative Impact, Auditability,* and *Documenting Trade-Offs* correspond to dos and don'ts of section 5 of our code, which deals with the "+1" step *Implementing Ethical Governance*.

In what follows, we provide a brief definition of each step and of the main issues we identified as deserving of ethical guidance. We shall do so also by presenting selected content from the code, drawing especially from the *key ethical ideas* of each section.

1) Data Acquisition and Generation

By definition, data-based services substantially rely on data. Thus, the very first step to consider is the collection of the data that is going to be the basis of the service. Data might be gathered from the data subject, e.g. by asking a client to provide it directly, or by recording or measuring user online or offline behaviour. For example, Facebook gathers data from its users, such as their posts, likes, their reactions or not-reactions on displayed commercials, etc. Similarly, an Internet of Things (IoT) platform gathers data by generating them directly from machine sensors. In addition, data (personal or anonymous) might be bought from a data provider. The result of this step is a raw data set, tagged with metadata. The key ethical ideas of this step are:

1. *You[1] have a duty to protect the informational self-determination of your clients and your employees and a duty to help them to make meaningful, autonomous decisions about their own privacy exposure.*
2. *You have to be transparent on which data you collect and what you are going to do with the data. Transparency includes being easily understandable.*
3. *To respect human autonomy, choices about design of user interfaces are as important as the content of legal contracts*[2]

2) Data Storage and Management

After data is collected, it has to be stored, and data management has to be defined and enforced. This includes prescriptions concerning secure authentication in order to ensure data privacy, and rules on updating or deleting data. However, there are more responsibilities associated with this step. Sometimes data is pre-processed. For example, data records might be filtered out and deleted because they appear to be wrong or felt to be not useful, for example because of missing data fields. Doing this has an impact on the following step in the data pipeline, i.e. *Data Analytics and Knowledge Accumulation*, because a filtered data set might lead to misleading conclusions in data analysis. For example, data lacking information about certain socio-demographic

[1] In the key ethical ideas presented in these notes, 'you' refers to a company developing data-based products or services.

groups may lead to models that are especially inaccurate when applied to those groups; measures taken for anonymization (e.g. replacing income figures with income brackets) may cause a data utility loss. Thus, a data analyst trying to derive any conclusions from a data set not only needs to know where the data comes from, but also what happened in terms of pre-processing before storing them into the database that forms the basis of the analysis. The key ethical ideas of this step are:

1. *Prevent harm deriving from violations of privacy and confidentiality and non-authorized usage.*
2. *Empower your clients to control their data.*
3. *Take responsibility for the quality of the data you use and manage, including their pedigree (such as restrictions of usage).*[2]

The second step results in a database management system containing the collected data, possibly in a preprocessed and augmented version.

3) Data Analytics and Knowledge Accumulation

The third step in the data pipeline consists of creating knowledge from data. This is the classical domain of the data scientists. Typical activities are

- performing statistical analysis on the data set, e.g. for revealing correlations between age and buying behaviour.
- using machine learning algorithms for finding relevant patterns in data or create predictive models, e.g. a model that assigns the probability of customers to terminate a phone contract in a given future timeline based on the socio-demographic and the usage data that is available.

The result of this step is either "knowledge" (e.g. "customers buying our products are typically below 30 years"), or a piece of software that processes input data into some output result. This includes automatic classification algorithms, detection algorithms, recommendation systems, and so on.

The ethical questions in this step are centred around the responsibility for creating, first, reasonably correct and, second, reasonably unbiased conclusions from data, as well as making clear on which assumptions the conclusions are made, what the area of application is, and why some biases are non-eliminable, while others were eliminated.

The first element (correctness or quality of conclusions) concerns the question "Are the inferences drawn from the data correct, and in which sense?" In technical terms, concepts like statistical significance, cross-validation or confidence intervals have to be applied to clarify the quality of the insight that is gained from the data. Moreover, reflection on the used assumptions and application area is even more critical. Often, models are built that are trained with data from one area, but then used in another application area. For example, face recognition systems are trained with databases that are often dominated by white male faces, leading to poor performance for black women [22]. Typically, the quality of the inference depends on the application area, and this has to be taken into account when building on such inference.

The relevant key ethical idea of the *Code* concerning reasonable accuracy, for this step of the data pipeline, says:

> *Data-based models may create harm by virtue of being inaccurate for a given application. Those who produce models are morally accountable for the proper communication of the limits of the knowledge derived from the data.*[2]

As the above example shows, a model having different predictive accuracy for different groups may be considered biased, and its use potentially discriminatory, especially if the unequal accuracy confers disadvantages to specific groups [23]. One key ethical idea in the *Code* pertinent to this step addresses the problem of bias:

> *Bias may have different sources. Some biases may create unfairness when the model is used in practice. It is the ethical responsibility of modelers, whenever possible, to specify the type and the extent of biases in the model, both bias that should be avoided and that cannot or should not be avoided, and act consequently. (Note: What counts as 'unfair' and as 'discrimination' is not defined univocally. Accept the existence of reasonable disagreement about the definition of 'unfair'. Be open to different perspectives and transparent about your own.)*[2]

This is one of the most complex claims contained in the document. It uses concepts (e.g. bias, unfair, discrimination) which are hard to define and which are thus more precisely defined in a glossary. It acknowledges that bias has different sources and different forms of discrimination may result from the deployment of models. A model may be said to be biased because it reflects structural social inequality affecting data generation (e.g. education or income levels reflecting gender or racial prejudices in educators or employers, or unequal opportunities), representation bias (when selecting the population, e.g. predominantly white and European), measurement bias (identifying which features and labels to use, especially when imperfect proxies are used) [24]. Some scholars, and many among media and justice activists tend to use 'bias' in a broader way, i.e. for any model that disproportionately assigns different outcomes to members of certain groups (disparate impact), irrespective of the reason for this unequal prediction. For example, some may call a model that displays ads according to a skewed (gendered) pattern biased, or discriminatory, even when this behaviour reflects the different (gendered) preferences of the users, or established (gendered) patterns of consumption in the population [25], [26]. In fact, in common or journalistic parlance virtually any model could be labelled biased when it reflects, produces or reinforces an inequality that society (or a part of society) considers morally objectionable. In the last years, the question of how algorithmic bias arises, and what can be done to avoid it, has become a major research area [27]. Even if many tools and knowledge have been developed for dealing with algorithmic bias, no consensus has been reached about the kind of bias that ought to be eliminated, as eliminating all biases is not a logically consistent goal in most circumstances [23]. It is the responsibility of the modellers to identify the biases they have an ethical duty to avoid, to check their models for bias, and to document bias that they cannot or should not avoid.

Summing up, a data scientist who creates an insight or a predictive model has to be careful to define in which areas the model is accurate, and how large the accuracy is, and he has to be transparent about this when making the model available to be used in a data-based service or product.

Furthermore, the code recommends that *any* kind of socially important bias embodied in the model should be made transparent, and unfair biases should be corrected.

4) ***Deployment of a data-based product or service***

The fourth step consists of using the insights that have been generated from data (either as "knowledge" and/or as model trained on data) as an element for a data-based service, and offering this service in the marketplace such that it can be used by clients, where clients may be individuals (such as users of digital B2C services) or companies/organizations.

It is only when the predictive models are deployed at the core of products or services reaching real users that some ethically problematic aspects of data collection and model training choices are actualized. It is in this step where, for example, people who should be treated equally are treated differently (for example caused by unequal accuracy or human discriminatory biases learned by the model). It is here that social trends of polarization, echo chambers, or the diffusion of fake news could be reinforced. It is here where people's behaviours is influenced – for the better or worse. The associated ethical issues are centred around the question of impact.

Our code distinguishes between two kinds of impact and two possible forms of harm: individual and collective. For the *individual* impact level the key ethical ideas are:

1. *The ethical use of data products requires assessing the impact of decisions based on data products on individuals and legal entities.*
2. *Possible harms of data-driven decisions involve: privacy violation, disadvantageous discrimination, loss of autonomy (including inability to challenge automated decisions), reputational harm, social or professional stigmatization, etc. Possible benefits include: lower prices, more pertinent recommendations, access to new products and services, etc. It is ethically mandatory that the benefits outweigh the harms, also in the long term. When harms and benefits affect distinct individuals, a question of justice arises. If your product treats some categories of individuals better than others, you should be able to justify this.*
3. *A valid reason for the unequal treatment of individuals is a justification that is relevant given the purpose of the service or model application. This will be different in different contexts, e.g. marketing, health, insurance, etc. A valid reason should be understandable by people who may ignore the technical details of how the model works.*[2]

For the collective level, the key ethical ideas are:

1. *Digital environments result from the combined decisions of different companies on their customers and have the potential to affect long-term societal trends.*
2. *Big data practices can contribute to good and bad social outcomes, by virtue of how companies interact with each other. Companies should highlight some of these outcomes, and every company should deliberate about its proper responsibility to avoid unintentional harm.*
3. *Digital environments that favour the development of intelligence, self-control, prudence, rationality, and openness to diversity positively affect the freedom and autonomy of users of these services. These human qualities are also important for a well-functioning, participatory democracy.*[2]

Of course, the *magnitude* of impact might depend on which and how many clients have access to the product or service. In a trial phase with a restricted number of customers, the potential of harm is typically much smaller than in a setting where a service is rolled out to a large public: Facebook cannot reinforce polarization if a substantial part of the population does not use it.

5) ***Implementing Ethical Governance***

The sequence of the four steps is not sufficient, however, because the goals singled out by distinct recommendations are interdependent. For example, product managers who deploy data-based services in the marketplace cannot avoid discriminatory decisions if they ignore the algorithmic bias in the algorithm. So, they depend on inputs from the modeller (step 3). Similarly, a modeller has to know where the data comes from in order to apply the models appropriately, and so depends on input from step 2. Similar dependencies obtain between companies that rely on each other products (e.g. data, models, etc.). This complexity generates a need for principles of ethical governance. These are provided by the following two key ethical ideas:

Data-ethical governance has two fundamental goals:

1) ensuring that data-ethical standards are respected within the company and properly aligned with the company's mission, values, and public image;

2) enabling ethical practices at the level of the data ecosystem (i.e. beyond the individual company). This can only be achieved if truthful information about the ethical standards of each company in the data ecosystem is available to other entities.[2]

VI. COMMENTS ON THE CODE

In the preceding section, we have outlined a structured way of describing ethical duties, using the data handling steps as a foundation. In this section, we add some additional comments.

A. Specialized companies in a data economy

We want to emphasize that the four steps may have different weight depending on the company; some may even not be present in a company. While larger companies may have in house capacity to perform all the data-related activities mentioned in this code of ethics, small companies may specialize only in one activity. For example, some companies may specialize in data collection and sell data to other companies; other companies may purchase data and generate statistical models based on this data; still other companies may purchase ready-made statistical models trained with data of other companies and use them to make decisions concerning their clients. Our *Code* can be applied for such companies as well. Care was taken to include all data-related business activities even of companies, which do not consider themselves as developing and offering data-based services. For example, brokering and trading with data might be the sole business activity of a company. In our 4-steps structure, this is located in step 1, as this is about collecting data and transferring them to someone else for further processing. Therefore, we explicitly envisage the case that a business actor is only active in a part of the four steps. The obligation to prevent data theft and misuse has been allocated to step 2, as this step covers all questions of storing and access control.

B. Assignment of responsibilities and ethical pedigree

The ethical quality of a data-based service or product depends on the activities of all steps. For example, if a data scientist creates knowledge by building a predictive model, using data that was acquired in a morally problematic way, the model may be considered ethically flawed. Ethical governance is the quality of an ecosystem which can only be enabled by ethical recommendations with a strong focus on procedures and proper communication. Agents in different points of the chain of responsibility are recommended to provide adequate information about the standards followed at each step of the data pipeline. Thus, a company (or department within a company) that relies on the product of some other company (or department), should seek information about the ethical pedigree of the previous steps. Similarly, a company (or department) that occupies an intermediate step in the data cycle should provide information about its own ethical standards to companies building on its results. For a large company, ethics requires organizing responsibilities to achieve ethical standards between different departments, and achieving an ethical pedigree should be a system-level goal for the organizational structures within the company. The responsibility to set these organizational structures in place is a higher-order managerial responsibility.

We do not assign ethical responsibilities to specific business roles, in the light of significant organizational differences among companies. In some firms, for example, the same business roles may supervise and be responsible for different steps. Thus, every business user of the guidelines is supposed to map the recommendations in the code to the appropriate organizational roles in the company. There are standard methods to implement ethics in organizations; in a follow-up project, we will analyse the adequacy of those methods to this Code.

C. Normative premises and values

The guidelines described here are prescriptive: their main content is a list of do's and don'ts, for each step in the data handling process. The satisfaction of these do's and don'ts is imagined to guarantee an ethically sound framework for the operations of a data-intensive company, from data collection to the generation of broad social impact on society. Thus, we conceptualize the overarching ethical commitment of any data-intensive firm as a commitment to act ethically in each specific step of the data pipeline and to enable the emergence of an "ethical pedigree" of data-driven products and services, within the company, and serving the informational needs of the information ecosystem. Ultimately, the measurable outcomes produced by implementing the Code are (ethical) process improvements. The Code contributes to them by helping individual decision-makers within the company to identify the impact of existing practices, the actions that ought to be taken, the responsibilities that need to be assigned, the new roles that need to be created, to fulfil the recommendations within distributed decision processes, and to support the organization-wide discussion.

As a source of inputs about the need for moral regulation and the relevant moral values, we have relied on the professional experience of the many group members who have been actively engaged in developing and deploying data-based services and products, as well as the expertise of some academic members with the themes and proposals discussed in the data ethics literature and in other ethical guidelines. Members of the Expert Group were also able to reflect on their experience as customers of many data-driven services. They have generated and criticized the guidelines with both perspectives in mind.

The ethical Code described in the paper has been written adopting the perspective of *ethical*, not *strategic CSR* [28], [9]. This means that the question asked in each meeting was how to achieve a fair balancing of the interests of multiple stakeholders, not what is the strategy that achieves a certain business goal (e.g. enhanced reputation). It does not follow, however, that the ethical guidelines are incompatible with the long-term profitability of the firm. If the attempt to capture a socially shared understanding of ethics has succeeded, implementing these guidelines aligns the company with customers' (reasonable) expectations. To mitigate the risk of reputation loss, a company must embrace and be able to respond to a range of considerations other than profit. In this sense, one could argue ethical CSR provides a safer route to long-term social sustainability, compared to anticipating reputational concerns strategically. Negative responses grounded in ethical values may be volatile and difficult to foresee. In choosing this approach, one trusts that the current understanding of the Expert Group, of reasonable and socially shared expectations, will illuminate the path. The second reason for choosing this approach is that the holistic identification of socially shared values and their implications for the practice of data science is less dependent on the specific context and constraint of a company. A strategic vision of ethical governance is, by contrast, highly context-dependent. Thus, the resulting Code is more widely applicable. For companies that consider business ethics mainly as a strategic tool to achieve specific business objectives, a more complex and personalized assessment for aligning strategy and ethics is needed, which these guidelines do not provide.

In summary, the guidelines collected in our Ethical Code aim to capture widespread ethical expectations from all stakeholders that might be affected by the data-based services under consideration. Thus, our recommendations are also relevant for companies whose commitment to 'ethics' is instrumental, e.g. that aim to reap the business benefits of being perceived as 'ethical' or that aim to avert government regulation. For this case, the Code would still be a reference frame for decision-making. For each recommendation, the company may assess the strategic risk of not following the recommendation, and ignore or reduce recommendations judged as not critical.

VII. CONCLUSION

In this paper, we provide an ethical, business-oriented, and methodological justification for the creation of a new set of ethical guidelines co-produced by companies and academics. We certainly do not advance the view that such guidelines, even if followed by the letter, would be *sufficient* to solve all the ethical problems emerging in the context of intensive data-based products or services. Yet we propose that a serious effort to implement these would reduce the probability of some violations of expectations concerning ethics. These, we believe, often occur due to the ignorance of practitioners of the ethical problems involved and the lack of sufficiently concrete guidance attuned to their needs.

REFERENCES

[1] https://data-service-alliance.ch/.

[2] Ethics Working Group - Swiss Alliance for Data-Intensive Services, "Ethical Code for Data-Based Value Creation V0.9." https://data-service-alliance.ch/innovation/ethics.

[3] V. H. Grassegger and M. Krogerus, "Ich habe nur gezeigt, dass es die Bombe gibt," Das Magazin, 03-Dec-2016. https://www.dasmagazin.ch/2016/12/03/ich-habe-nur-gezeigt-dass-es-die-bombe-gibt/

[4] Independent High-Level Expert Group On Artificial Intelligence Set Up By The European Commission, "Ethics guidelines for trustworthy AI," European Commission - Digital Single Market, Brussels, Apr. 2019. https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai

[5] Regulation on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation 2016/ 679). 2016.

[6] H. C. Wallich and J. J. McGowan, "Stockholder interest and the corporation's role in social policy," in A new rationale for corporate social policy, W.J. Baumol, R. Likert, H.C. Wallich and J.J. McGowan, Eds. New York: Committee for Economic Development, 1970, pp. 39–59.

[7] A. B. Carroll, "A Three-Dimensional Conceptual Model of Corporate Performance," AMR, vol. 4, no. 4, Oct. 1979, pp. 497–505.

[8] A. B. Carroll, "The pyramid of corporate social responsibility: Toward the moral management of organizational stakeholders," Business Horizons, vol. 34, no. 4, Jul. 1991, pp. 39–48.

[9] M.-D. P. Lee, "A review of the theories of corporate social responsibility: Its evolutionary path and the road ahead," International Journal of Management Reviews, vol. 10, no. 1, 2008, pp. 53–73.

[10] FAT ML, "Principles for Accountable Algorithms and a Social Impact Statement for Algorithms," no date. [Online]. Available: http://www.fatml.org/resources/principles-for-accountable-algorithms.

[11] The Open Data Institute, "The Data Ethics Canvas," 05-Aug-2017. https://theodi.org/article/data-ethics-canvas/

[12] Department of Health & Social Care - UK Government, "Initial code of conduct for data-driven health and care technology," *GOV.UK*, 05-Sep-2018. https://www.gov.uk/government/publications/code-of-conduct-for-data-driven-health-and-care-technology/initial-code-of-conduct-for-data-driven-health-and-care-technology.

[13] J. Manske and T. Knobloch, "Datenpolitik jenseits von Datenschutz." Stiftung Neue Verantvortung, Oct-2017.

[14] MyData.org, "MyData Declaration." https://mydata.org/declaration/ [Accessed: 2019-Jan-29 15:10:39].

[15] Université de Montreal, "Montreal Delcaration for a Responsible development of AI," 2017. https://www.montrealdeclaration-responsibleai.com/the-declaration.

[16] Future of Life Institute, "AI Principles," 2017. https://futureoflife.org/ai-principles/ [Accessed: 29-Jan-2019].

[17] Data & Society, "Governing Artificial Intelligence," https://datasociety.net/output/governing-artificial-intelligence/

[18] AI4People, "Ethical Framework for a Good AI Society," Atomium European Institute for Science, Media, and Democracy, Brussels, 2019. http://www.eismd.eu/wp-content/uploads/2019/03/AI4People%E2%80%99s-Ethical-Framework-for-a-Good-AI-Society.pdf

[19] C. Villani, "For A Meaningful Artificial Intelligence: Towards A French And European Strategy," Parlamentary Mission assigned by the Prime Minister of France Édouard Philippe, Mar. 2018. https://www.aiforhumanity.fr/pdfs/MissionVillani_Report_ENG-VF.pdf

[20] P. Chapman, J. Clinton, R. Kerber, T. Khabaza, C. Shearer, and R. Wirth, "CRISP-DM 1.0." 1999. ftp://ftp.software.ibm.com/software/analytics/spss/support/Modeler/Documentation/14/UserManual/CRISP-DM.pdf

[21] A. I. R. L. Azevedo and M. F. Santos, "KDD, SEMMA and CRISP-DM: a parallel overview," in *IADIS Proceedings of Informatics 2008 and Data Mining 2008*, H. Weghorn and A. P. Abraham, Eds. Amsterdam: IADIS, 2008, pp. 182–185.

[22] J. Buolamwini and T. Gebru, "Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification," in *Conference on Fairness, Accountability and Transparency*, 2018, pp. 77–91.

[23] S. Mitchell, E. Potash, and S. Barocas, "Prediction-Based Decisions and Fairness: A Catalogue of Choices, Assumptions, and Definitions," *arXiv:1811.07867 [stat]*, Nov. 2018.

[24] H. Suresh and J. V. Guttag, "A Framework for Understanding Unintended Consequences of Machine Learning," *arXiv:1901.10002 [cs, stat]*, Jan. 2019.

[25] M. Ali, P. Sapiezynski, M. Bogen, A. Korolova, A. Mislove, and A. Rieke, "Discrimination through optimization: How Facebook's ad delivery can lead to skewed outcomes," *arXiv:1904.02095 [cs]*, Apr. 2019.

[26] S. Biddle, "Facebook's Ad Algorithm Is a Race and Gender Stereotyping Machine, New Study Suggests," *The Intercept*, 04-Apr-2019.

[27] A. Chouldechova, "Fair prediction with disparate impact: A study of bias in recidivism prediction instruments," *arXiv:1610.07524 [cs, stat]*, Oct. 2016.

[28] D. Windsor, "Corporate Social Responsibility: Three Key Approaches," *Journal of Management Studies*, vol. 43, no. 1, 2006, pp. 93–114.