Swiss Alliance for
Data-Intensive Services

**Expert group Data Ethics**

**Version 3, June 10 2019**

**For Public Consultation**

# Ethical Codex for Data-Based Value Creation

## Table of Content

# Introducing the Codex

## Purpose of this Codex and addressees

This Codex is a guideline for all organizations that offer services or products based on data. The primary focus are businesses, in particular Small and Medium Enterprises (SMEs), which do not have specialized in-house expertise for the ethical aspects of data usage. For those, the Codex will be helpful for developing services that are consistent with the ethical expectations of customers, employees, or society. However, the Codex may also be used by any other organization dealing with data.

The purpose of the Codex is to address the ethical questions that arise when using data, and to suggest best practices in the form of “do’s” and “don’t”s.

Ethical decisions around data usage are coupled to the question of *how and to what end* data is being used. These questions arise along the four steps of

- data acquisition,
- data storage and access control,
- data processing and knowledge generation, and
- usage of data-created knowledge in a concrete context.

The formal structure of this Codex is following these steps and is addressing the corresponding issues of each step. So, the Codex creates a guideline for dealing with data for all steps of developing, deploying and maintaining data based services.

## Why respecting an Ethical Codex for Data-Based Value Creation?

Prudent companies will strive to reduce the risk for bad reputation that derives from violating the ethical expectations of customers, employees, or society. Adequately dealing with ethical tensions is part of sound business practice, sustainable profit making, and sustainable growth. GDPR and other legal frameworks *alone* do not provide sufficient guidance; acceptance by stakeholders and the possible reputation damage potential provide a further benchmark to act on. If these risks are not considered, trust might be lost and a major reputation damage might be provoked.

In addition, individuals working in a company usually are attracted to do the good, for its own sake, and strive to pursue it to the extent that the business circumstances permit it. There-

fore, ethically compliant companies can also improve their reputation as employers, which helps them to attract talents that care for the ethical conduct of their employer.

Finally, ethical behavior may create business opportunities and competitive advantage in a market of consumers with increasing ethical expectations, in particular regarding privacy and discrimination issues. Following the recommendations of this Codex helps to build trust and reputation in the marketplace, which is a competitive advantage in modern service economy.

## Who should use this Codex?

This document can be used as a reference framework by management and individuals in every organization dealing with data. Especially in larger organizations, the decisions with respect to data usage are being made in many different parts of the organization. For example, a corporate IT department is responsible for storing and data access control, data scientists are responsible for data-based knowledge generation, and product managers are in charge of developing and deploying data-based services in the marketplace. In other organizations, the responsibilities might not be as clearly attributed. In each case, however, an ethical data usage requires that all data-related decisions (end-to-end) are coordinated and fit together, based on a common set of data-related corporate values.

In a world that is working increasingly in eco-systems with a plethora of partners, activities are not only distributed among different divisions within a single company, but also more and more among different companies. For example, the step “knowledge generation from data” might be outsourced to a data-science company that uses provided data, possibly enriches it with self-provided data, and returns a predictive model. To illustrate another example, the step “data storage and access control” might be outsourced to a cloud service.

In all instances, however, the management of each firm has to deliver a data governance structure to ensure that the handling of data (respectively information) meets ethical standards end-to-end. This governance structure must address ethical issues emerging at different stages of the data life cycle and the relation between them. It must also consider that the ethical standard of a larger ecosystem emerges from the collaboration of the actors. That is, the ecosystem as a whole can only be ethical if each actor is transparent enough for the other actors: For example, a company that uses input from another company must be able to assess the trustworthiness and ethical standards of these inputs (e.g. data, models, etc.).

In the context of our four-step model for data-based business, an input for a specific step comes with an “ethical footprint”, depending on the data-handling processes of the prior steps, and this footprint becomes a part of the “ethical footprint” of the output of this step. For example, if a data scientist builds a predictive model using data that has not been authorized

to be used for such a purpose, the resulting model is unethical, not because the data scientist ignored ethical guidelines for modeling, but because the input data was flawed. It should be part of the responsibilities of each step to analyze the “ethical footprint” of its input and clearly communicate the “ethical footprint” of the output, based on the input and the activities of this step. The output has to be transferred to the next step together with this “ethical footprint”.

The Codex helps individual decision makers to position their specific field of responsibility within the full context, to understand how their decisions might impact ethical questions of other decision makers, to generate reasonable interfaces between distributed decision processes, and to support the organization-wide discussion.

## How to use the Codex?

Often, ethical considerations lead to ambiguities and conflicts of values. In ethical decision-making, conflicting values have to be weighed against each other. Depending on this weighing, the decision might be different. In these situations, the Codex should help to clarify what the ethical dilemma is, and which values are touched. The Codex strives to support the decision maker to take a decision that is well founded, consistent with the corporate values, and transparent, such that it can be defended both internally and externally, in particular in discussions with customers and society.

This Codex gives recommendations in the form of “do”s and “don’t”s. As all ethical decisions have to be taken in a specific context, the way how exactly to follow these recommendations may vary and needs to be adapted to the circumstances of the individual situation.

Endorsing this Codex means to do the following three things:

1. Assess all “do”s and “don’t”s for the applicable steps,
2. Take a conscious decision as to whether or not to follow the recommendations, to which degree, and in which concrete form,
3. Document these decisions in a written form and give justifications for them

Thus, for each deployed service, the way how the data is used from data acquisition until leveraging data-based knowledge in a specific form in the context of the service is both transparent and justified, and can be defended against internal and external criticism.

# The four steps of data usage

Data-based services allow an organization to generate value based on data. For offering such services, data has to be acquired, to be stored, to be processed, and the result of this processing has to be integrated in a well specified service offer.

The following four steps indicated these necessary activities:

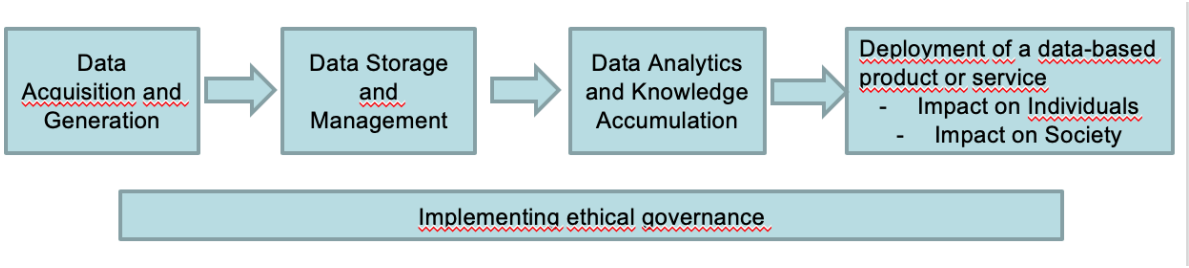


In the following, these four steps are described in detail, in particular their inputs and their outputs.

**Data acquisition:**

- This activity consists of acquiring the data that is being used for offering a service in the end. Data (including meta-data) may be acquired by directly interacting with individuals (e.g. letting them provide personal data on a website), from sensors/machines (IoT), by purchasing them from data brokers, or by other means. Note that data is not only raw data, but may also be data that has been generated by processing other data (e.g. an age distribution of a population is data that can be acquired. This data has been generated by someone else who processed a large number of age information of individuals).
- The ***output*** of this step is a set of digitized data that may be enriched with meta-data.

**Data storage and management:**

- The ***input*** of this step is the digitized data of the previous data acquisition step.
- The activity of this step consists of storing this data, defining and implementing rules of how to update or delete data sets, defining and implementing rules of giving access to this data. In addition, any pre-processing of the acquired data prior to storing is part of this step (e.g. data cleansing such as remove duplicates or obviously wrong data, invalid entries, etc.), enhancing the data with other data, linking data with other data, eliminating data records from the originally acquired data set for whatever reason).
- The ***output*** of this step is a database that is used for subsequent data analysis.

**Data analytics and knowledge accumulation:**

- The ***input*** of this step is a database.
- The typical activity is the generation of knowledge from the data. This might be done by using descriptive statistical methods (e.g. calculation of histograms, empirical distributions, etc.). It may also be done by building a predictive model with machine learning (ML) methods (including statistical learning, neural networks, etc.).
- The ***output*** of this step is derived knowledge either in the form of data that is coding the knowledge (e.g. in the form of correlation coefficients, means or quantiles, etc.), models (e.g. a trained neural network), or documents describing what information can be derived from the data.

**Deployment of a data-based product or service:**

- The ***input*** of this step is knowledge artifacts, as generated in the preceding step.
- The activity of this step is offering a data product, e.g. a data-based service, or changing any real-world processes based on the derived knowledge (e.g. selection of a specific marketing message based on the knowledge about the addressed person). Typical activities are:
    - differential treatment of individuals, e.g. giving or denying access to resources (e.g. loans), or individual pricing,
    - giving access to opportunities (e.g. job opportunities, both in HR, and in job ads),
    - giving access to information (personalized information),
    - influencing individual behavior in particular directions (recommendations, personalized information).

  This step comprises all activities that have an impact in the real world (as opposed to just generating knowledge) by using the data-based knowledge in a specific implementation. We may differentiate
    - Impact on individuals (persons, legal persons, …): This includes positive and negative impact that is produced because the data-based knowledge is used. An example for negative impact is unwanted discrimination (resulting from “algorithmic bias”).
    - Impact on society: this includes large-scale impact on the structure of society. An example is a fragmentation of society by generation of information filter bubbles due to personalized information delivery systems.
- The **output** of this step is a fully implemented data product which can be used by individuals, organizations or institutions.

# The role of values and the link between ethical principles and prescriptions

The moral prescriptions (dos and don'ts) for each step of data usage are preceded by 3 key ethical ideas, that differ from step to step. These ethical ideas are mid-level ethical propositions half way between practical directions (the dos and don'ts) and the abstract ethical values typically discussed by moral philosophers. Existing ethical guidelines concerning big data and AI mention moral principles and values at a high level of abstraction, for example:[1]

1. Human rights
2. Well-being
3. Control and autonomy, including of information and (digital) identity
4. Transparency, explainability
5. Accountability and auditability
6. Responsibility
7. Privacy and intimacy (protection)
8. Liberty (or freedom)
9. Shared benefit/prosperity, common good, fairness, equity, justice, equality and non-discrimination
10. Solidarity
11. Democratic participation
12. Sustainability

The Codex incorporates these values/principles. All these principles inspire the dos and don'ts below, but with an emphasis on the context in which they have to be operationalized. Sometimes the application of the value/principle to a specific application domain (in our case one of the four steps of the data life cycle) requires a somewhat longer explanation of the implications. We believe that this is more useful for the intended audience of this document. The star sign (*) near a value term means that we provide a general definition of the value in our **glossary.**

Human rights are not mentioned explicitly here. However, the key ethical ideas are strictly related to the human rights that are typically mentioned as important for big data and artificial intelligence. This concerns in particular:

1. Respect for the right to *privacy* is promoted by recommendations in ch. 1, 2, 4, and 5
2. Respect for the right of avoiding discrimination is promoted in recommendations in ch. 3, 4 and 5.
3. Respect for the right to political participation and freedom of expression is promoted by the recommendations in chapter 5.

[1] The list is not complete, and sometimes there are overlaps between the principles/values.

Finally, some of the dos and don'ts can also be interpreted in terms of human virtues. Several accounts of virtues such as honesty, justice, modesty and prudence, or vices, such as vanity or greed, emphasize the importance of stable dispositions, the ability to reciprocate benefits, the awareness of limitations and the capacity to be transparent about it, the constant struggle to avoid self-serving bias and understand the potential benefits and risk of what one does on progressively larger groups of people.

# The Codex

## 1. Data Acquisition and Generation

**Key ethical ideas:**

1. You have a duty to protect the informational self-determination* of personal data providers such as 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* about which data you collect and what you are going to do with the data. Transparency includes being easily understandable.
3. To respect the autonomy* of personal data providers, web design choices affecting user experience are as important as the content of legal contracts, especially if they nudge data subjects to release more data.

**Do's**

1. DATA TRANSPARENCY Find simple ways to explain which data you collect, to explain your privacy policy and to explain how you will use your clients' data (which purpose, which results). Use summaries and images to increase readability. Be creative in how you communicate this information, e.g. by using storytelling approaches, graphs and illustrations. Inform your clients if you combine data your clients give you voluntarily with other data about them (e.g. acquired from data brokers).
2. DATA CONTROL Provide your clients online tools to make choices about how their data will be used. If you are build a 'nudge' into the design of your service, it should always be in the clear interest of the client, clearly documented in advance, and you should mention it already when collecting the data.
3. PROFILE TRANSPARENCY If you use clients' data for building profiles about them, your clients should know. Explain your profiling with examples that the average client can understand.
4. RECIPROCITY Engage your clients if you want to get more data from them than the minimum strictly necessary for providing your service. If clients give you their informed consent to collect unnecessary data, they should obtain adequate value (not necessarily, money) in return. Better still, enable them to demand, obtain and reuse their data.

**Don'ts**

1. OPACITY Do not use long and complex terms and conditions without summaries and other ways to signal important information. Collaborate with other companies to generate new standards for more readable, understandable, and classifiable terms and references.

2. HYPOCRISY Do not try to hide how you derive economic value from data, be upfront about it.
3. SELF-SERVING NUDGING Do not use web design and user experience to nudge your clients to give you access to more data than strictly necessary, exploiting their inertia, lack of attention and other biases.
4. NON-TRUSTWORTHY ACQUISITION Do not purchase data from partners who are not trustworthy and transparent about their data collection practices and any restrictions coupled to the usage of the data. Your corporation's standards for data acquisition should be guaranteed by any external data provider from which you acquire data.
5. DATA GREED Big data does not necessarily imply big value: refrain from acquiring data you do not know how to use, yet.

## 2. Data storage and management

**Key ethical ideas:**

1. Prevent harm deriving from violations of privacy* and confidentiality* and non-authorized usage of data stored in your IT-system.
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).

**Do's**

1. SOURCE TRACEABILITY When you store acquired data, add metadata to them that specify the source and what kind of usage is allowed for this data for which period. Choose a level of granularity that is adequate to the purpose. In big data approaches, the purpose can have a broad scope, but still it needs to be declared. Only store data that is tagged with this information.
2. TRUSTWORTHINESS OF SOURCES Assess the trustworthiness of your data providers. Keep track of where the data you use comes from, in particular when you combine data sets. Take care that usage restrictions are passed on if data sources are combined, and that all "parent" data sets are documented.
3. ACCOUNTABILITY Define clearly the rules for access control (who, when, under which conditions, tracking of access, …). Define rules for deleting data sets. Implement appropriate cybersecurity measures which are coherently applied to all IT service models, e.g. cloud-based or hybrid. Document any change and adaptation of data management requirements when transitioning to different models.
4. PRIVACY PROTECTION Adopt a variety of privacy-protecting techniques based on an appropriate assessment of the privacy risk. Be technologically up to date.

5. TRANSPARENCY ABOUT THE DATA LIFETIME If it is necessary to store data for a very long time (e.g. for accountability, or service lifetime reasons), explain this to your clients at the time of data collection. Foresee what to do with data of inactive accounts. If possible, give your customers discretion and provide non-complicated options for customers to restrict data storage.

**Don'ts**

1. UNLIMITED COLLECTION Do not store personal data that is not necessary for providing, improving or expanding your service. This is the best protection of privacy and against liabilities.
2. POOR ANONYMIZATION Anonymization does not simply mean deleting your customers' names. Make sure that anonymized data is truly anonymized, i.e. cannot be traced back to the original person it relates to. Simply deleting obviously sensitive fields is often not enough, as the remaining data can be used to reconstruct the deleted data field.
3. UNLIMITED DURATION Critically assess if the planned lifetime of data is reasonable given the goals of your service.

## 3. Data analytics and knowledge accumulation

**Key ethical ideas:**

1. Data-based models may create harm by virtue of being inaccurate. Those who produce models are morally accountable for the proper communication of the restrictions and the limits of the knowledge that is being derived from the data.
2. Data-based models may have bias*. Some biases may create unfairness when the model is used in practice. It is the ethical responsibility of modelers to specify the type and the extent of biases in the model, whenever possible, both bias that should be avoided and that should not and cannot 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.

**Do's:**

1. EXPLAIN METRICS Explain and justify the performance metrics used for modelling. Check for sampling bias, indirect and direct discrimination, check model selection, proper train-test setup. Data sets used for training models should enable the assessment and, if possible, the mitigation of biases resulting from poor training data.
2. AVOID UNFAIR DISCRIMINATION Declare what you consider to be "eliminable bias". Describe and declare the kind of bias that is still in the model (and the justification of it), if any. Develop in-house expertise on problems of eliminable discrimination

or unfairness in machine learning, including the technical (statistics, computer science) methodologies. Document your activities to eliminate discrimination and bias in the model. You should document not only your choice of technique or fairness goal but also the reasoning leading to it. For example, determine if the data reproduces an unfair status quo that should not be preserved or reinforced.

3. PROMOTE INTELLIGIBILITY Create explanations of how your model works and out its outcomes relate to data input, and test them with non-tech people and users of the model. At least when all other things are equal, prefer models that can be more easily explained to users of your model.
4. EXPLAIN LIMITATIONS Be honest with users about the limitations of your models, especially if lacking awareness of these limitations increases the risk of controversial uses.
5. ASSESS AND LIMIT HARMFUL USES Assess possible harmful uses of your models and describe them, for creating awareness among users of the models. Identify ways to make it more difficult to use your models in socially harmful ways.

**Don’ts**

1. ONE-DIMENSIONAL METRICS Do not ignore ethical considerations of privacy, unfairness and explainability, even if they have a cost in terms of accuracy of profit optimization. When building models, explore the trade-off between accuracy/profit optimization and those other values.
2. VANITY Do not hide limitations and problematic uses of the data-based knowledge (e.g. a predictive model) when discussing its adoption with potential users of this knowledge.
3. ETHICS NEUTRALITY Do not assume that developing a model or derive knowledge from data is ethically neutral, just because you are not in charge of using the model on people. Take ownership of ethical choices affecting modelling.
4. DO NOT PROMOTE BLIND TRUST Do not advise people who will use your model/knowledge to blindly trust a model they do not understand, but rather promote the conscious use of any models based on a reasonable understanding of their operation and logic only.

# 4. Deployment of a data-based product or service

## a. Impact on individuals

**Key ethical ideas:**

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: service improvement, 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 run. When harms and benefits affect individuals unequally, 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.

**Do's:**

1. EVALUATE METRICS BEFORE USE When using a data-based model within the context of a service offering, make sure that you have a clear understanding of the performance metrics used for modelling, and of the kind of bias that is still in the model (and the justification of it), if any. If purchasing a model from a third party, obtain information whether some fairness metrics, as specified in the technical literature, has been considered and implemented when training the model. If the model changes over time (e.g. because of feedback loops, or in self-learning systems), this evaluations have to be repeated after every change and in case of self-learning in a regular manner, e.g. monthly.
2. FAIR TREATMENT OF INDIVIDUALS If the algorithm you use treats different customers differently (e.g. limiting access to a service or discount), you should be able to identify and explain a valid reason for it. Thus, prefer using a model that is easier to interpret and explain, at least when all other things are equal.
3. UNDERSTAND LIMITATIONS Understand the limitations of your models, and make sure that the model is used in a way that delivers acceptable performance (and avoids unfair bias) in your specific case.
4. ENABLE DISPUTABILITY Develop awareness of how inaccurate or biased decision may harm clients. When appropriate, provide procedures that enable clients to challenge outcomes based on models, if harmful to them, for example procedures allowing contestation and re-examination of some decisions.

**Don'ts**

1. EXPECTATION BETRAYAL Do not use analytics that reveal aspects of your clients that would violate the expectations of a typical client. For example, inferences about sexual orientation, mental health, and social status from behavioral data may violate these expectations, at least in certain contexts (this is should be assessed by someone with an understanding of the culture in which the product is deployed).
2. ILLEGITIMATE MANIPULATION Do not use analytics with the aim of manipulating users based on their identified weaknesses (knowledge that results from analyzing their data). If, e.g., nudges in the choice architecture are used, provide a moral justification for them (e.g., they may be necessary to protect the users' security or privacy).
3. AD HOC DEVIATIONS When you have a policy that relies on a data-based model to make decisions, do not deviate from the decisions/suggestions of the model in an ad hoc manner. Ad hoc adaptations may cause more ethical problems than redress them. Better to develop a strategy for dealing with predictions or classifications that appear counterintuitive and to report problems deriving from the faulty applications of models.

## b. Impact on society

**Key ethical ideas:**

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, also beyond affecting the interests of individual users of services for which companies are strictly responsible. Companies can highlight some of these outcomes, and every company should deliberate about its proper responsibility to avoid unintentional harm.
3. Digital environments that favor 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.*

**Do's**

1. ENGAGE WITH STAKEHOLDERS Provide contacts of persons with responsibility for ethics in your company, which may be contacted by individuals or organizations worried about the impact of a data-based product on society. Identify relevant ethical issues together with stakeholders and prepare an impact analysis on the issues.
2. PROMOTE DIGITAL WELL-BEING Assess the impact of your data-driven services on desirable human traits, such as health, intelligence, self-control, prudence, rationality and openness to diversity. Seek opportunities to promote those valuable condi-

tions and traits through your products and initiatives associated with them. Avoid services that undermine valuable human qualities.

3. ASSESS PRIVACY SPILLOVERS Assess the impact of your data-driven services on norms concerning privacy in society. Are your products making it difficult for individuals or specific groups of individuals to hide aspects of themselves or their lives that they may reasonably want to hide from you or others? Seek opportunities to contribute to a society that values personal privacy in all its aspects, for instance by enabling privacy-conscious options for your customers.
4. CONSIDER POSSIBLE HARM BEYOND YOUR CLIENTS Be aware of how the data you collect or output of your models can be misused by not so well-meaning parties. Be also aware of the risk associated with models of risk that may unintentionally stigmatize or threaten the reputation of entire groups of individuals. Take reasonable precautions to mitigate such risks.

**Don'ts**

1. MISUSE IGNORANCE Do not use data and models without examining how they might be misused.
2. UNCONTROLLED PUBLICATION Do not make models public without reasonable precautions, if they have a high chance of being misused.
3. HARM NEGLIGENCE Do not ignore harm that may be produced on individuals whose personal data you do not own or process (e.g. members of groups whose reputation may be affected by models).

## 5. Setting up ethical governance

**Key ethical ideas:**

Data-ethical governance has two fundamental goals:

1) Ensuring that ethical standards are respected *within* the company in all steps of data handling, and are properly aligned with the company's mission, values, and public image, by making sure that all decisions about how data is being used are synchronized and lead to ethically justifiable data-based business end-to-end;
2) Enabling ethical practices at the level of the data ecosystem (i.e., beyond the individual company), by enabling each actor of the ecosystem to create a good ethical standard for his own activities. This can only be achieved if truthful information about the ethical standards of each company in the data ecosystem is available to other entities.

**Do's:**

1. SELF-ASSESSMENT Companies using data for their business should ethically handle all the relevant steps of the data usage process. They should also report on their

ethical standards and practices to other companies which use their outcomes, and the public. This information can be produced by identifying the chapters of this Codex which are relevant to the own core business activities and the performance indicators of the relevant do's and don'ts.

2. TRACEABILITY Generate and demand information for sustaining ethical data practices:
   2a) If your product depends on data and/or models of other companies, additionally obtain all information necessary to use those data and/or models in an ethical way. For example, if you use a model from someone else for your own product, obtain clear information about the limits of this model.
   2b) If your product depends on data and/or models of other companies, make sure that what you use has been produced in an ethical way. For example, if you use a model of someone else, make sure that it has been trained with legitimate data.
   2c) If other companies build on your data/models, provide them with the information they need for steps 2a and 2b.
3. STRUCTURE Each company should identify individuals and roles to ensure that ethical standards are respected and to determine accountability for wrongs. More precisely, it should identify a) management roles responsible for defining and setting up ethical standards and goals in your company, b) specific duties of every role in the company which is involved in respecting data-ethical standards and improving (to the extent feasible) the governance processes.
4. PROCESSES Implement processes to ensure that ethical standards are met. The reliability of these processes should be tested and measured in practice, not just assumed. In particular, make sure that all necessary information is passed from one step of the data pipeline to the next one. Make sure that the final product satisfies ethical standards end-to-end. Processes should be tested and improved over time.
5. PRIORITIES Identify the relation between the values in this code and the company principles that are relevant for your company. Determine the data ethics priorities for your company which result from your commitment to this code and align them with your mission as a company.

**Don'ts**

1. IGNORANCE OF THIRD PARTY STANDARDS Do not ignore that your ethical standards are strongly influenced by the ethical standards of other companies you interact with.
2. SCAPEGOATING Top management should not identify scapegoats, that is, should not deny corporate responsibility, when it is the case that responsibility is more broadly shared.
3. DISENTANGLE POWER AND RESPONSIBILITY Do not distribute responsibilities unfairly: more power should always correspond to higher accountability.

# Glossary

EDITRORIAL REMARK: The Glossary is not yet complete. Reviewers are invited to suggest additional terms that should be included in the Glossary

| | |
|---|---|
| ALGORITHM | An algorithm is 'any well-defined computational procedure that take some value, or set of values, as input and procedures some value, or set of values, as output' (Cormen et al. 2001).<br>When the word 'algorithm' is used in this codex, we consider as 'algorithms' those human artifacts stemming from the training of machine learning models on digital data, in order to generate predictions to assist or automate decision-making. Therefore, we consider a special class of algorithms, i.e. those *resulting* by solving a given machine-learning problem. A machine learning problem 'can be precisely defined as the problem of improving some measure of performance P when executing some task T, through some type of training experience E' (Mitchell, 1997). Training experience E is represented by (digital) input data, which are preprocessed and formatted for the machine-learning problem under consideration. Among the most relevant tasks T, one comprises classification (i.e. the specification of a label for each data point, from a finite set of possibilities) and regression (i.e. the computation of a numerical value of interest for any data point). For simplicity, we do not discuss reinforcement-learning problems (Sutton & Barto 1998). |
| AUTONOMY | |
| BIAS | The term 'biased data' collapses retrospective injustice (societal bias) with concerns about non-representative sampling and measurement error (statistical bias)' (Mitchell et al 2018). Societal bias is for example the fact that income or education scores may reflect unequal opportunities, or the prejudices of those giving the scores. It is the gap between the world as it should and could be and the real world. Statistical bias is the gap between the sample and the world. |
| CONFIDENTIALITY | |
| CONTROL | |

| | |
|---|---|
| DEMOCRACY | |
| DISCRIMINATION | |
| FAIRNESS | |
| FREEDOM | |
| GDPR | General Data Protection Regulation of the European Union Regulation (EU) 2016/679 |
| ML | machine learning: |
| PRIVACY | |
| RESPONSIBILITY | |
| SELF-DETERMINATION | |
| SME | Small and Medium Enterprise |
| TRANSPARENCY | |
| | |
| | |