The CoDE
Code of Data Ethics

Edited by Enrico Panai

Version 1.0

FORHUMANITY

DISCLAIMER

The philosophy, tools and methods presented in this paper are often a practical elaboration of the works published by Luciano Floridi. If in case some sentence has been quoted without the correct form or reference, please excuse me and report it. This paper only tries to organise some thoughts to enable an operational approach to data ethics. While the credit for the content therefore lies with the original authors, any misinterpretation of the cited work is my complete responsibility.

NOTE

Defined terms are **bolded** throughout the document.

CONTRIBUTORS

Gamze Büşra Kaya, Ryan Carrier, Zoe Mariangela Cocchiaro, Divya Dwivedi, Marianna Ganapini, Kathrin Grosse, Inbal Karo, Rohan Light, Lukas Madl, Luca Nannini, Sundar Narayanan, Cheryl Rego, Carolina Sanchez, Maria Santacaterina, Emmanuelle Shaaravi, Alessio Tartaro, Paolo Volpe.

I thank all the people who contributed to this document (listed alphabetically). I remain the only person responsible for the shortcomings.

GRAPHIC DESIGN

Gianpiero Nughedu

STRUCTURE

The Code of Data Ethics (CoDE) consists of two parts, a PUBLIC PART on general principles and methodology and a PRIVATE PART, where the values and approaches used at each level of the organisation and at each stage of the pipeline are explained.

PRINT AND ELECTRONIC DISTRIBUTION RIGHTS

PUBLIC PART

Table of Contents

Why we need a code

Introduction

Chapter 1

The Code of Data Ethics is adopted to assist members of the organisation in making ethical choices related to data, data processes, and AI, algorithmic and autonomous systems.
This document introduces the data ethics framework to define the organisational boundaries, the structure of the Code of Data Ethics, and the roles of a Data Ethicist (i.e., Chief Data Ethics Officer or a Data Ethics Committee) responsible for the maintenance and enforcement of this code.
The fundamental principle that inspired this document is based on *The Ethics of Information* (Floridi, 2010) according to which the main concern of a digital ethicist "should be the ecological management and wellbeing of the infosphere". This approach clearly shows the environmental approach to the ethics of data, information and artificial intelligence on which this code is built.

What is a Code of Data Ethics?

A Code of Data Ethics (CoDE) is a set of principles, definitions, procedures, methods, and practices that helps develop ethical reasoning on issues involving data, information, algorithms, and digital infrastructure. The Code of Data Ethics provides the enterprise with the tools to align the values expressed in its Code of Ethics with practice. The CoDE uses a level of informational abstraction in order to achieve the following goals: i) to shape governance, ii) perform ethical analyses of systems (Axiological Analysis), iii) predict possible ethical risks (Ethical Analysis Forecasts), and iv) allocate responsibility (distributed morality) for actions generated by complex agents (Multi Agent System) in the space to which Data Ethics applies. In a nutshell, the Code of Ethics identifies business values and principles, while the Code of Data Ethics provides the theoretical framework and operational tools to be able to govern digital processes, from data to autonomous artificial intelligence systems.
An ethical framework on data, AI, algorithmic or autonomous systems provides useful perspectives which are suitable for fostering moral evaluations in line with the ethical principles expressed by the organisation in its Code of Ethics.
This framework is based on the principles of information ethics (Floridi, 2013b), which is concerned with the well-being and prosperity of the digital environment in which the organisation operates.

What is a Code of Data Ethics for and how is it applied?

A Code of Data Ethics serves to drive the Digital Governance of an organisation by

- helping to shape the processes and methods used by data stewards and data custodians;
- defining strategies for allocating moral responsibility for the consequences of handling data and information;
- applying the principles expressed in the code of ethics (by developing satisfactory axiology to solve problems related to data, information, algorithms, infrastructures, and practices) ;
- highlighting the limits of its implementation;
- using ethics as an innovation trigger and analysis tool to avoid opportunity costs by enhancing ethical reasoning

What is it important for?

An informational approach to data-related moral evaluations should no longer be postponed. The influence of the digital in relating to the world (through its epistemological impact, its capacity to reontologise[1] reality, and the autonomy of agents from an ethical point of view) imposes a change in how we allocate moral responsibilities and how we make corrections to models (both technical systems, such as mere algorithms, and socio-technical, such as the organisation itself) to mitigate negatively morally-loaded outcomes. In fact, we may argue that ''big'' morally-loaded actions can be the result of many, ''small'' morally-neutral or morally-negligible interactions among agents constituting a multiagent system (MAS), which might be human, artificial, or hybrid. The agency capability of AI systems makes this approach particularly important in order to redistribute moral responsibility, to develop risk mitigation mechanisms and to avoid opportunity costs due to the apprehension to use an AI system.

In an informational environmental there are various pressure points "where a difference can be made to good and evil" (Floridi, 2016) based on the nature of agents (virtue ethics), on the nature of actions (deontology and consequentialism) and on the status of patients (understood as the recipients of the action). Classical ethics, even when analysing the consequences of an action, focuses on the agent and their faculty of agency. In contrast, information ethics, inspired by medical ethics, takes care of the patient and for these purposes it analyses the entire moral situation. According to Floridi (2016), "the shift in perspective is from an agent-oriented ethics, which cares about the individual development, social welfare and ultimate salvation, to a patient-oriented ethics, which cares about the affected system's well-being and ultimate flourishing. With an analogy, the ethical discourse may focus on the cook, on the cooking or on the cooked" (Floridi, 2016).

How is it applied?

The ethical framework is relevant to whomever is responsible for the quality of moral assessments in order to 1) create an environment conducive to ethical reasoning, 2) make explicit the tools needed for identifying the elements that constitute an ethical situation, for defining the levels of the threshold of moral relevance or negligibility, 3) establish the governance that enables the implementation of the code itself.

Who is concerned by it?

The framework provides the conceptual elements for action. It is not directly necessary for everyone, but it is essential for an appropriate use of digital ethics in an organisation. Regardless of the appellations proposed in this document, those responsible for the data ethics must have a *well established competence* in digital ethics: because doing data ethics without a solid framework would be like being a doctor without a solid competence in anatomy.

AI, or at least its massive dissemination and application, is relatively recent. Its great potential has attracted many students, engineers and researchers. At the same time, serious concerns about the abuse or misuse of AI have fuelled the debate. Some of the concerns are real, while others are fuelled by a misunderstanding of the technology. However, they have prompted the spread of people who are concerned, often rightly so, about the possible negative consequences. AI does indeed impose a major paradigm shift, but to articulate it, it must be governed with the right tools. Data, AI and digital ethics requires theoretical knowledge and practical skills to prevent or mitigate possible ethical risks and to not miss opportunities for development due to poorly formulated concerns. Indeed, opportunity costs can be the hidden face of ethics if its practice is pursued by agents who are not sufficiently trained. At present, there are only a few specific courses in the market, and some are still very approximate. However, it is possible, and desirable, that a standardisation of the competences of data, AI and digital ethicists sees the light soon.

1 In the meaning of designing and constructing anew the very nature of the infosphere, that is, of the environment itself, of the agents embedded in it and of their interactions.

Where should it be used?

Data ethics is a latent space in companies that operate in informationally mature societies (Panai forthcoming). This means that ethical choices are already being made about data, but that the ethical decision-making process is not explicit: it is not supported by any axiological analysis. In order to prevent errors and mitigate ethical risks, the solution is to make the role of data ethics explicit within the organisation by:

- Identifying a person who is responsible for ethical choices about data and its interface with Artificial Intelligence (AI) systems: the Chief Data Ethics Officer (CDEO).
- Supporting the Chief Data Ethics Officer with a Data Ethics Committee (DEC) in order to mitigate personal biases and to assist them in the process of making ethical choices.
- Creating and maintaining a specific Code of Data Ethics (CoDE), which sets out the principles and core values in data management, indicates the chain of responsibility and regulates the use of data and algorithms as much as possible.

When should it be used?

The society in which we live is already a mature information society, so the use of a data ethics framework is already necessary. However, one cannot imagine applying it without an adequate and explicit business ethics framework already being defined in the organisation's Code of Ethics. This implies a value priority, not a temporal one between the two frameworks. Given that information ethics enables the development of efficient models for business ethics (Floridi, 2013, 277-291), one can invite organisations lagging behind in the creation of an ethical infrastructure to engage in parallel work (Code of Ethics & Code of Data Ethics).

Why is there a need for a Code of Data Ethics?

In any mature information society, every organisation should have a Code of Data Ethics. However, it becomes indispensable when the organisation handles large amounts of personal and sensitive data or when it uses systems with evolutionary dynamics, such as AI, algorithmic, and autonomous systems. In particular, the peculiar nature of the informational agents (where the agency is not only human, but can be also artificial or hybrid) makes the allocation of moral responsibility more problematic and not solvable through the publication of internal codes of conduct (as it would be where the agency is only human). A Code of Data Ethics is necessary because the agent involved in the moral situation is no longer just a human agent or an artificial agent triggered by human logic (as in deterministic systems), but can be an autonomous agent triggered by information obtained by inferential processes.

Who is affected by a Code of Data Ethics?

In mature information societies, a Code of Data Ethics shapes the decision-making process of the whole organisation. However, it is especially addressed to those who deal with data governance (data stewards and data custodians) and those who directly manipulate and wrangle data (data workers, data scientists, developers).

The Chief Data Ethics Officer (CDEO) and the Data Ethics Committee (DEC) should be responsible for drafting and maintaining the CoDE. They must have recognised skills in data ethics and must have been directly endorsed by the organisation's board. In the case of organisations that are particularly advanced in the use of AI systems, the Chief Data Ethics Officer (CDEO) may be joined by an AI ethicist with sufficient statistical skills: the Chief AI Ethics Officer (CAIEO).

When should a Code of Data Ethics be created?

A Code of Data Ethics needs to draw on the principles of the organisation in order to support axiological analyses and moral assessments of digital governance and compliance with digital regulation. Therefore, it is necessary that the Code of Data Ethics is preceded or accompanied by a sufficient or mature Code of Ethics.
In the absence of a code of ethics or a set of principles that are not clearly shared or defined, this code may be inspired by the Unified Framework of Ethical Principles for AI (Floridi et al., 2018):

- Beneficence: understood as "promoting well-being, preserving dignity, and sustaining the planet";
- Nonmaleficence: understood as "operating 'within secure constraints'", preventing "of infringements on personal privacy", enhancing "cautions against the various negative consequences of overusing or misusing AI technologies";
- Autonomy: understood as "the power to decide to decide", in its three form (decide, to decide-to-delegate, decide to decide again);
- Justice: understood as "promoting prosperity, preserving solidarity, avoiding unfairness";
- Explicability: understood as "incorporating both the epistemological sense of intelligibility—as an answer to the question 'how does it work?'—and in the ethical sense of accountability—as an answer to the question 'who is responsible for the way it works?'".

The framework consists of four principles from the bioethical tradition (Beneficence, Non-maleficence, Autonomy and Justice) plus an additional principle required by the fact that AI is a new form of agency (Explicability), and it played a valuable role in other documents: AI4People (Floridi et al. 2018), Ethics Guidelines for Trustworthy AI published by the European Commission's High-Level Expert Group on AI (HLEGAI 18 December 2018, 8 April 2019); the OECD's Recommendation of the Council on Artificial Intelligence (OECD 2019); The EU AI Act (in progress); Beijing AI Principles (Beijing Academy of Artificial Intelligence 2019); Rome Call for an AI Ethics (Pontifical Academy for Life 2020)

Where does digital ethics fit into digital governance?

From an operational perspective, data ethics is used as an applied or non-standard ethics and is positioned under the umbrella of the enterprise's shared business ethics framework. However, the purely technical aspect of data ethics allows its representative Data Ethicists (Chief Data Ethics Officer and the Data Ethics Committee) to interact directly with those responsible for digital compliance (ex. Data Protection Officer for GDPR in the European community and in the UK), Cybersecurity (Chief Information Security Officer) and data workers. In the informational environment of the organisation, data ethics is used to enforce the principles and uphold the shared values of the organisation.

What is its structure?

The Code of Data Ethics (CoDE) consists of two parts, a *Public Part* on general principles and methodology and a *Private Part*, where the values and approaches used at each level of the organisation and each stage of the pipeline are explained. The *Private Part* concerns the operational approach to data ethics. Although it will soon be published as an explanatory supplement to the *Public Part* on, the *Private Part* should remain an operational document and thus distributable within the organisation that is to use it. The *Private Part* is necessary to guarantee the ethical operability of the processes.
After this introductory chapter, in Chapter 2, we briefly present the ethical framework. In Chapter 3, we introduce practices concerning the responsibilities and liabilities of people and organisations. In Chapter 4, we approach problems posed by the collection and analysis of large. In Chapter 5, we address issues raised by the increasing complexity and autonomy of algorithms.

The foundations of this code

Ethical Framework

Chapter 2

Data Ethics is the branch of ethics that studies and evaluates moral problems related to data (including generation, collecting, recording, curation, processing, dissemination, sharing, and use)[1], algorithms (including AI, artificial agents, machine learning, and robots), and corresponding practices (including responsible innovation, programming, hacking, and professional codes), in order to formulate and support morally good solutions (i.e., right conducts or right values) (Floridi & Taddeo, 2016)

For a broader understanding, it may however also be useful to know the definitions used by the Open Data Institute (ODI)[2] and the United Nations Office for the Coordination of Humanitarian Affairs (OCHA)[3]

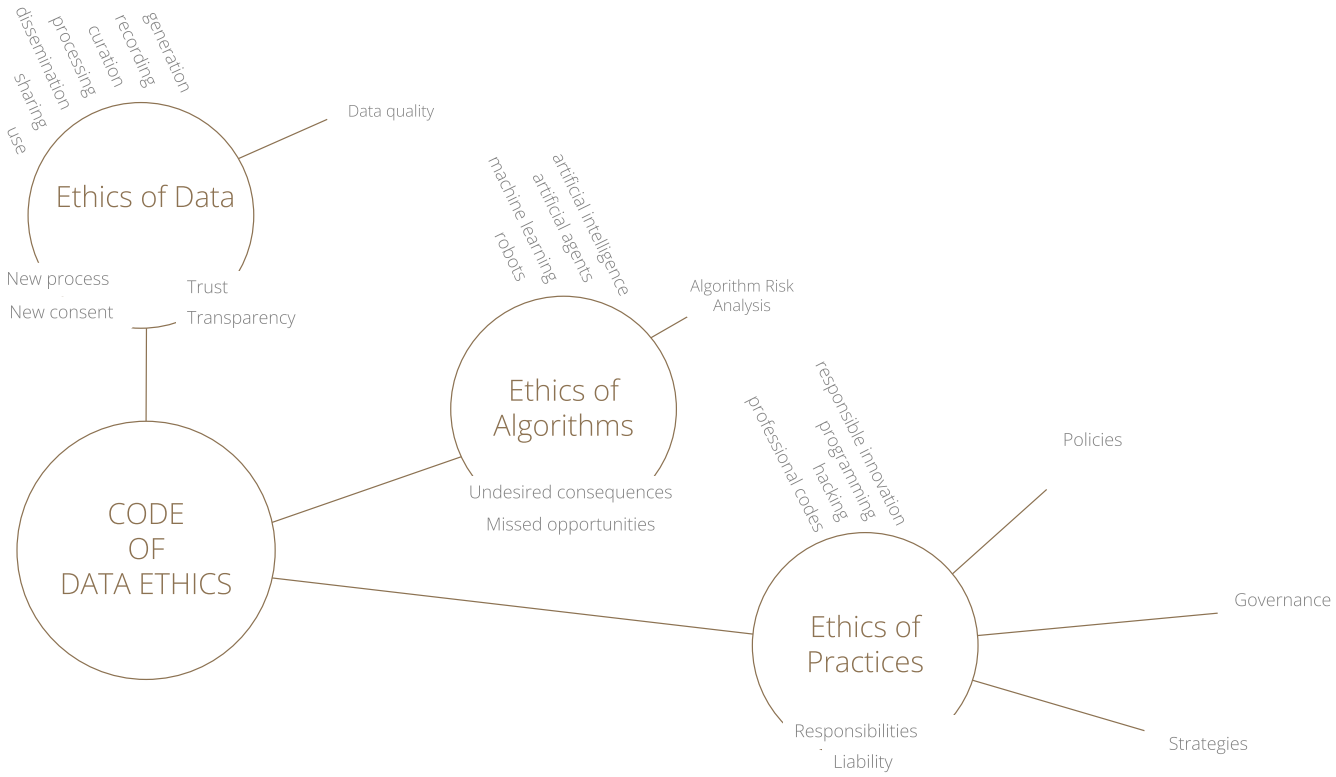

1 More specifically in the algorithmic approach "collection, labeling, cleaning, transformation and reduction, training, testing/validation"

2 Open Data Institute (ODI): "A branch of ethics that evaluates data practices with the potential to adversely impact on people and society – in data collection, sharing and use"

3 United Nations Office for the Coordination of Humanitarian Affairs (OCHA): "The branch of ethics that studies and evaluates moral problems and offers normative guidance related to data, algorithms, and corresponding practices. Common ethical issues in data management include issues of fairness, validity, bias, ossification, transparency and explainability, anonymity, privacy, and ownership of data and insights."

According to Floridi & Taddeo (2016), Data Ethics consists of three macro areas (as depicted in figure 1). In this CoDE, the order of use is functional, and the interests of each macro-area have been aligned with business needs

- ethics of practices,
- ethics of data,
- ethics of algorithms.

The three areas represent the set of ethics applied to data processing. The framework on which Data Ethics is based has its roots in information ethics (Floridi, 2013), which provides the theories, methods and tools to apply these ethics.

ELEMENTS OF THE ETHICAL FRAMEWORK

The world can be represented in many different ways, but in order to make decisions we must first choose the perspective we want to use. Choosing the right perspective, through the **method of levels of abstractio**n, "plays an absolutely crucial role in how we handle any information process, and so in how we negotiate our interactions with the world" (Floridi, 2013, p. 30). Specifically, the ethical framework of this code is built on an **informational level of abstraction**, implying that the world is composed of **informational entities** that have minimal intrinsic moral value and thus may deserve to be respected. Therefore, in order to deal with a universe of informational entities within the organisation we must use an ethics adapted to a world where information may be processed from data by other pieces of information, the algorithms. Briefly, if we live in a digital world and we cope with informational entities we should use a digital ethic. In particular, this code uses an instance of information ethics, the Data Ethics developed by Floridi & Taddeo (2016).
Data Ethics shapes the digital governance of the organisation by acting at different levels of granularity involving three ethical macro areas: ethics of practices, ethics of data and ethics of algorithms.
The set of these macro-areas is used to build an ethical infrastructure within the organisation (also called **internal infraethics**). In particular, the ethics of practices is necessary for:

1) supporting the organisation in complying with regulations
2) aligning practices with the set of principles and values defined by the organisation, generally in the ethical charter or in the Code of Ethics
3) defining internal policies, roles and accountability
4) avoiding **ethical deviations**
5) helping management to choose what is socially preferable, going, "over and above the existing regulation" (Floridi, 2018), in other words using a post compliance ethical approach (defined "**soft-ethics**[1]")

In the informational world (or infosphere), some informational entities possess **agency**, namely the ability to take an action that produces a particular effect. Thanks to their **agency**, informational entities can become elements of a moral situation, either as emissaries of the action (agents) or as receivers of the action (**patients**). Understanding a **moral situation** is fundamental to understanding how an action

1 Great examples of the use of soft ethics are visible with the EU GDPR and the UK Children's Code

propagates. However, compared to the moral actions used by classical ethics, the level of informational abstraction forces us to model the moral situation (perhaps by using very convenient framework as the Object-Oriented Programming) with the aim of being able to 1) understand the propagation of the moral action, 2) allocate the moral responsibilities of the actions.
The **propagation of moral action** allows the identification of key milestones in a digital system lifecycle that need moral evaluation. In the case of deterministic systems it is mainly the ethics of data that provides the tools for moral evaluation, while when the algorithms are based on inferential models then the ethics of algorithms provides the right tools, such as the **map of algorithms**. It is often assumed that algorithms are used (1) for transforming data into evidence and (2) trigger and motivate action. But these operations, when performed by (semi-)autonomous algorithms such as ML, (3) complicate the allocation of responsibility. Consequently, the ethics of algorithms deals with concerns related to epistemic factors (*inconclusive*, *inscrutable*, and *misguided evidence*), normative factors (*unfair outcomes* and *transformative effects*), and one factor that is both both epistemic and normative (*traceability*) (Tsamados et al., 2022).
This raises the need for an appropriate body to decide the necessity and the proportionality of the use of AI algorithms. This observation opens the way to the ethics of practices, which attempts to answer some crucial questions such as “Who is responsible, accountable, reliable, culpable? Who is qualified to handle AI algorithms and autonomous systems? Who will guarantee the quality of ethical reasoning independently?”.

Allocating moral responsibility is fundamental to improve processes and create a trustworthy environment, however systems involving AI, algorithmic and autonomous systems cannot be considered as monolithic agents, but as elements of a complex system, also called **Multi-Agent Systems** (MAS). In this case the allocation is more complex because a set of morally neutral actions can cause morally-loaded consequences. So we need the **distributed morality** approach that applies a recursive **back propagation** for allocating punishments and rewards and **common knowledge** for increasing the chances to prevent ethical risks.
All the moral evaluations that are done from governance to AI systems, may always be satisfied by the **Chief of Data Ethics Officer** and the **Data Ethics Committee** according to the chosen **axiological analysis** of the organisation.
Briefly, the person or committee responsible for the data ethics of an organisation must have the tools and power to promote the flourishing of the informational ecosystem of the organisation and the entire infosphere.

LEVELS OF ABSTRACTION

The method of levels of abstraction (LoA) makes it possible to choose a good question answer. Already used in mathematics (where it is closely related to model theory), in computer science (mainly in object-oriented programming or OOP), LoA has been used by information philosophy to shift the focus from humans to informational entities. This shift makes it possible to build tools that are better suited to the agency capacity of artificial intelligence systems. In simple words, a precise level of abstraction allows one to take a point of view and develop one’s decision-making process within that perspective. It should not be considered as a relativistic approach, but as the tool that allows an analysis to be satisfied depending on the goals pursued.

What LoA for?

The method of levels of abstraction allows a point of observation to be chosen. As the level of observation changes, so do the "observables". In practice, specific observables refer to a specific level of abstraction (LoA).
In short, choosing the right level of abstraction allows one to formulate the right questions for the problem, and thus find a solution.

How it works

There are many different ways to travel. If we decide to travel, we can choose to do so using a motorbike or a sailing boat (or any other means of transport). The choice of the means of transportation will depend on the initial goal. Do we want to travel on land or at sea? Do we want to travel as a family, as a couple or alone? Do we want to travel with little or no pollution? Once we have chosen our goal, we position ourselves on a level of abstraction. For example, if we choose to travel by sea in a sailing boat, we will need to obtain nautical charts and know the tides. If, on the other hand, we travel by motorbike, we will need road maps. Following the initial idea of the journey, application is only possible by choosing the right **level of abstraction (LoA)**.
Once the LoA has been chosen (let's say travelling at sea with a sailing boat), we know that we will have to take nautical charts with us, but which ones? There are indeed charts with different scales. However, if we need to enter a port in Brittany, we won't be able to use a chart showing the entire Atlantic Ocean. We will need a more detailed chart, with indications of ports, depths, colour of lighthouses, etc. Without this granularity of detail, we will not be able to navigate in Brittany. Without this granularity of detail, we risk running aground on a rock. So we can only achieve our goals by choosing the right **level of granularity (LoG)**. For the Atlantic ocean, a map with little detail, and for entering a port in Brittany, a higher level of finesse.
When Data Ethicists have to make a moral assessment, they must first understand the LoA and then the LoG. This allows them to take action on the digital governance of the organisation and the quality of information in a process.
There are many different ways of travelling. The motorbike or the sailing boat are just two ways of interfacing with travel. We must therefore be aware that, when analysing a phenomenon, we are faced with several possible **gradients of abstraction (GoA)**. This diversity requires us to include multidisciplinary and multicultural approaches in our processes to avoid debasing the informational richness of less conventional abstractions.

Who can use it?

In principle, everyone should be trained to use the LoA method to address a discussion, because it limits misunderstandings while maintaining an appropriate level of complexity.
"An adage in human-computer interaction known as the law of conservation of complexity or Tesler's Law, states that for every system there is a certain level of complexity that cannot be reduced. In other words, every application has a certain degree of complexity that the developer or user has to deal with (Saffer, 2010, p. 136.). We might suggest using the same law for this paper or in general for addressing ethical issues.
The risk of reducing ethical complexity too much in the principles of a general ethical framework on data and AI may correspond to an increase in the complexity of implementing that framework. In short, we must be aware that we cannot have simple solutions for complex issues" (Panai and Light, forthcoming).
Furthermore, as a tool also used by developers or software designers, it provides a shared communication platform for the development of ethical solutions.
It therefore provides the Data Ethicist (Chief Data Ethics Officer or Data Ethics Committee) with a fundamental tool for communicating with developers.

When should it be used?

We should use the LoA method at the beginning of each kind of analysis (from ethical to technical) because... “the idea of a ‘level of abstraction’ plays an absolutely crucial role in how we handle any information process, and so in how we negotiate our interactions with the world, and therefore in how we develop our philosophy of information, including our information ethics” (Floridi, 2013, pg 30).

Why is it so important for the CoDE?

The method of levels of abstraction is important for a Code of Data Ethics for two reasons: 1) it serves to choose the initial level of abstraction, the informational one; 2) it is useful to have the ability to frame the moral situation and make the right assessments at different levels of granularity.

What does it mean that we use an informational level of abstraction?

All this code is based on a precise level of abstraction, the informational one. According to this level of abstraction, the world is interpreted as an environment consisting of informational entities (some of which have agency capabilities), their properties, interactions, processes and relations. The world could have been interpreted through a level of informational abstraction even before the digital revolution, but it is precisely the new binary technologies of information and communication that make the use of this level of abstraction indispensable. Digital technologies need to ‘envelop’ the world in order to function. In this sense, data is no longer inside the computer, but computers and digital devices work better because the world is made of data. In the informational world, the model that interprets the data is … data. In other words, the program that opens a (spreadsheet) file ... is also an (executive) file. Choosing the informational LoA allows information ethics to provide an intrinsic value to informational entities. Information ethics “does not refer to the moral value of well-formed and meaningful data (such as an email). What the ethics of information suggests is that we adopt an informational approach (technically, a level of abstraction) to the analysis of Being [ed. in philosophy, being is the material or immaterial existence of a thing.] in terms of a minimal common ontology, whereby human beings as well as animals, plants, artefacts (and hence emails), and so forth are all interpreted, insofar as they are entities, as informational entities” (EoI Floridi, 308).

This non-anthropocentric approach allows for a change in perspective. So, “instead of trying to stop agents treating human beings as informational entities, we should rather ask them to realise that, when they deal with personal and private information, they are dealing with human beings themselves, and should therefore exercise the same care and show the same ethical respect they would exercise and show when dealing with other people, living bodies or environmental elements” (EoI Floridi, 259).

The informational LoA also renews the classical model of the moral situation, no longer composed only of human agents, but also of artificial agents or groups of them, the multiagent systems (MAS).

THE MORAL SITUATION

The **moral situation** is the fundamental object of ethics. It can be defined as any action taken by an agent and which has an impact on a recipient. We often talk about **moral dilemmas**, but these are much rarer and involve situations in which a difficult choice must be made between two or more alternatives that are equally undesirable. Most commonly, the moral situation involves actions involving preferable alternatives: to print or not to print a document; to choose between a more or less environmentally friendly server; to put a product on the market that is not in line with the organisation's Code of Ethics; etc.

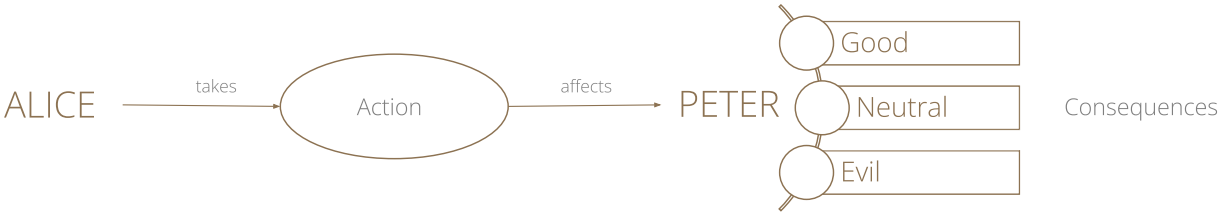

Thus every action has moral consequences, however some consequences remain '**morally negligeable**' and thus do not need ethical intervention, while others may be '**morally significant**'. In this case the moral charge may be positive (in which case praise and reward will be distributed) or negative (in which case blame and punishment will be distributed, and mechanisms will need to be implemented to mitigate the ethical risks).

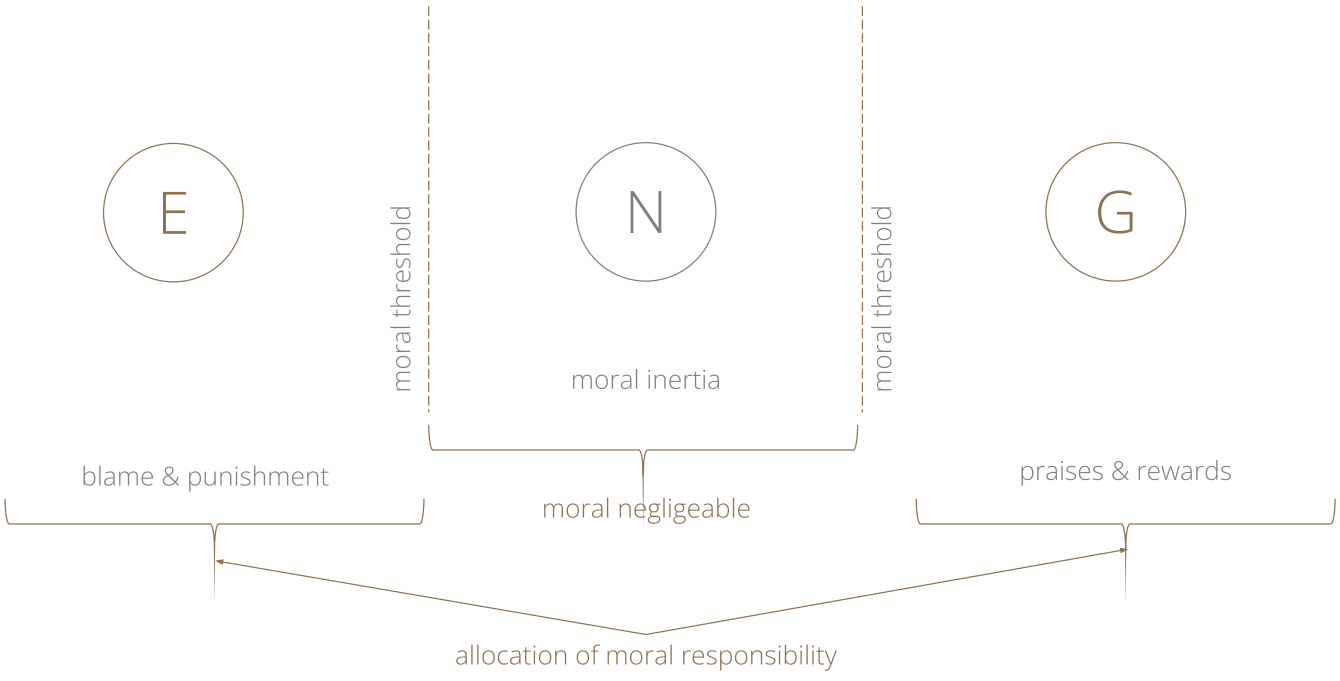

The role of an ethicist is to create an environment conducive to generating positive actions within moral situations. In this sense goodness must be understood as the absence of evil, thus including neutrality. In practice, an ethicist works for actions to be neutral or good. Due to the extension of entity classes generated by artificial agents, the strength of an ethicist lies in their ability to understand and assess the moral charge of an action (Floridi, 2013).

Let us take the example of printing a document. The action of a single individual may be considered “negligible”, but the sum of all documents printed by a company may have an ecological impact[1] . On the other hand, if the organisation has developed mechanisms to reduce printing when not necessary and taken steps to offset its carbon footprint by investing in reforestation, then not only is it avoiding a negative impact, but it is also taking a positive action.
In practice, the ethicist uses all available information to reduce the uncertainty of the consequences of an action. In addition to the law that proscriptively determines what not to do (can’t), the Data Ethicist analyses the moral situation in order to understand how to position the **moral threshold** that indicates when the consequences of an action fall into the area of **moral inertia** and are therefore morally negligible or when they assume a positive or negative moral charge. This allows the Data Ethicist to intervene on possible prevention or remedy mechanisms.

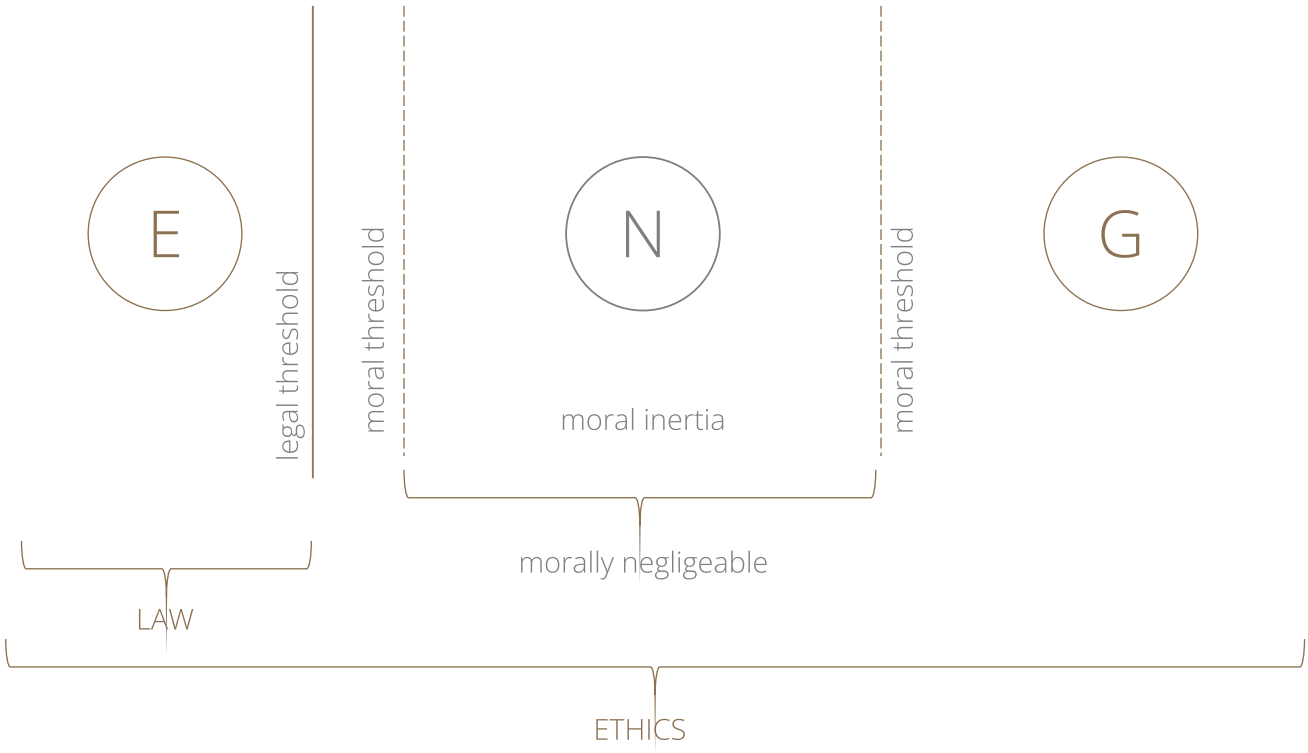

1 Also known as “many hands problem” (van de Poel, 2015) or “the tragedy of the commons” (Hardin, 1968, 1998)

According to classical macroethics

Back to moral action. In classical macro ethics, the moral situation consisted of an agent (Alice) and a receiver (Peter) who are human. Roughly, we can say that the classical macro ethics focused on one of the elements of the moral situation: virtue ethics on the quality of the agent, Kantian-based deontology[1] on the quality of the action, and consequentialist ethics on the quality of the consequences on the recipients.

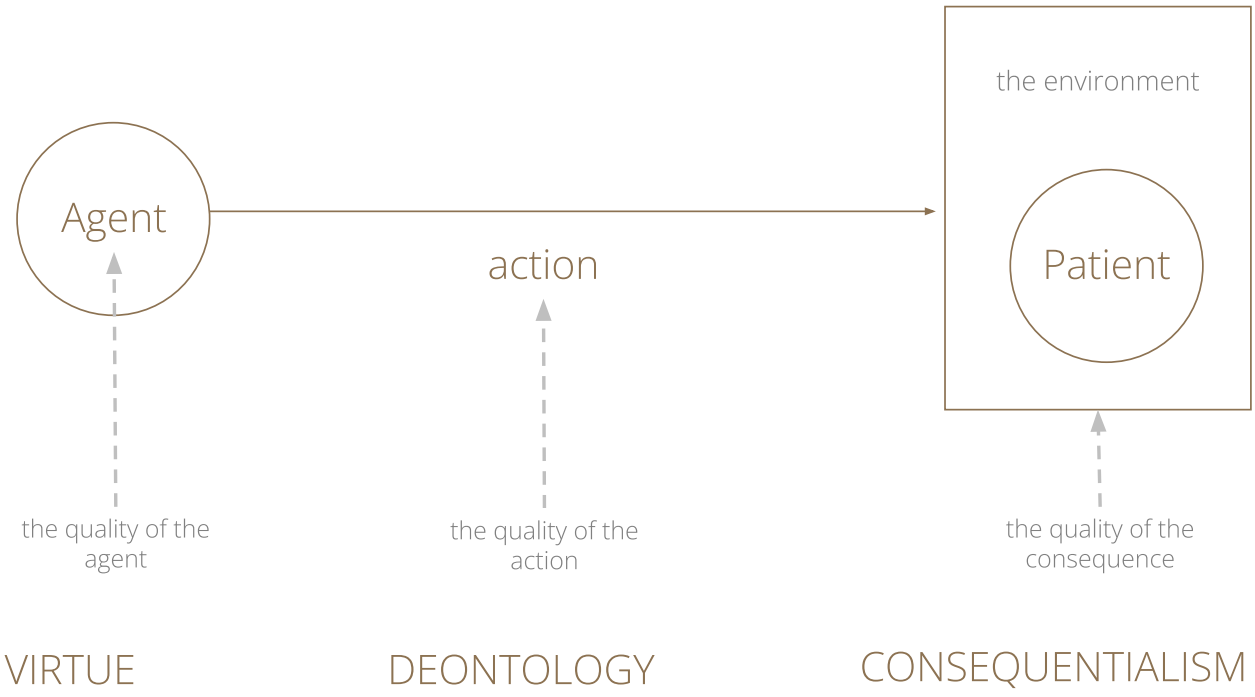

According to information ethics

Information ethics changes the level of abstraction compared to classical macroethics and considers all elements as **informational entities**, emphasising the importance of the moral situation as an **envelope** within which agents (human, artificial or hybrid) perform actions that impact on recipients or patients (human, artificial, hybrid or passive informational entities). This paradigm shift makes it possible to address moral situations generated or mediated by information technologies. To better understand how a moral situation develops, we can draw inspiration from the Object-Oriented Programming (OOP) approach to identify all the elements of a moral situation. In computer science, object-oriented programming (OOP) is a programming paradigm that allows the definition of software objects capable of interacting with one another through the exchange of messages. Particularly suitable in contexts in which interdependent relationships can be defined between the concepts to be modelled (containment, use, specialisation), a field that more than others succeed in exploiting the advantages of object-oriented programming is that of graphic interfaces. “The OOP approach provides a very flexible and powerful methodology with which to clarify and make the concept of ‘informational object’ precise as an entity constituted by a bundle of properties” (Floridi, 2013 p. 104). In this way we can have a practical application to moral valuations in an information environment.

1 A whole theoretical chapter could be written on the misuse of the term ‘deontology’ in some professional ethical deontologies.

An **agent (1)** activates an **information process (3)** thanks to the information it has at its disposal. This information is only a personal representation of the world and constitutes the agent's information **shell (4)**. The agent will act according to the **factual information (5)** concerning the specific moral situation in which it is involved. The information process influences a **patient (2)** and the relative personal world of information, the **shell (4)**. This is the **envelope (6)** within which the moral situation develops, which in turn is located in the information environment surrounding us, the **infosphere (7)**. Now since all these elements are informational elements, an action (or information process) has an impact on the whole informational ecosystem.

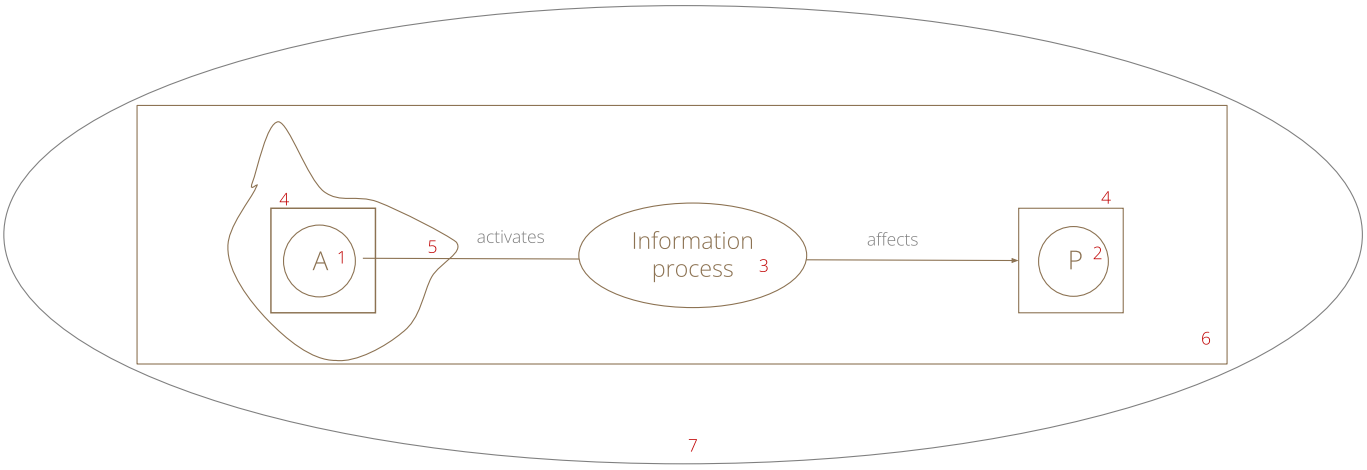

1) A = the moral agent
2) P = the moral patient
3) M = the moral action, as an interactive information process

4) *shell* as A personal world information
5) *factual information* = information about the moral situation
6) *envelope* = the moral situation
7) *infosphere* = the general environment

Floridi, L. (2013). *The Ethics of Information*. Oxford University Press. Chapter 6 p.103-133

Thanks to this interpretation, it becomes easier to understand why a dataset may have biases. Imagine a group of male developers constituting the dataset (i.e. the shell) with which the agent (which could be a human being, a software or an artificial intelligence system) will activate an informational process. There is a risk that the shell built by men alone will be embedded with the personal worldview of those men. This leads to the question of when to correct such distortions. One should not have diversity because it is politically correct, but because it is an ethical risk mitigation tool. Having a group of developers who prize diversity means allowing the shell to be as consistent as possible with the factual information available. At the same time, diversity can be considered as an innovation tool, allowing for more inclusive design and thus greater market impact. For as Albert Camus wrote, "the evil that is in the world almost always comes of ignorance, and good intentions may do as much harm as malevolence if they lack understanding" (Albert Camus, The Plague, trans. Stuart Gilbert (New York, NY: Vintage), p. 131).

The moral situation is to be imagined as the envelope within which an agent acts. Parallel to the concept of envelope used in robotics to describe the range of motion that a mechanical arm has, the concept of envelope in ethics must be conceived as the range of action of an agent. Determining such a space is complex and requires a good understanding of the ethical space involved by an action.

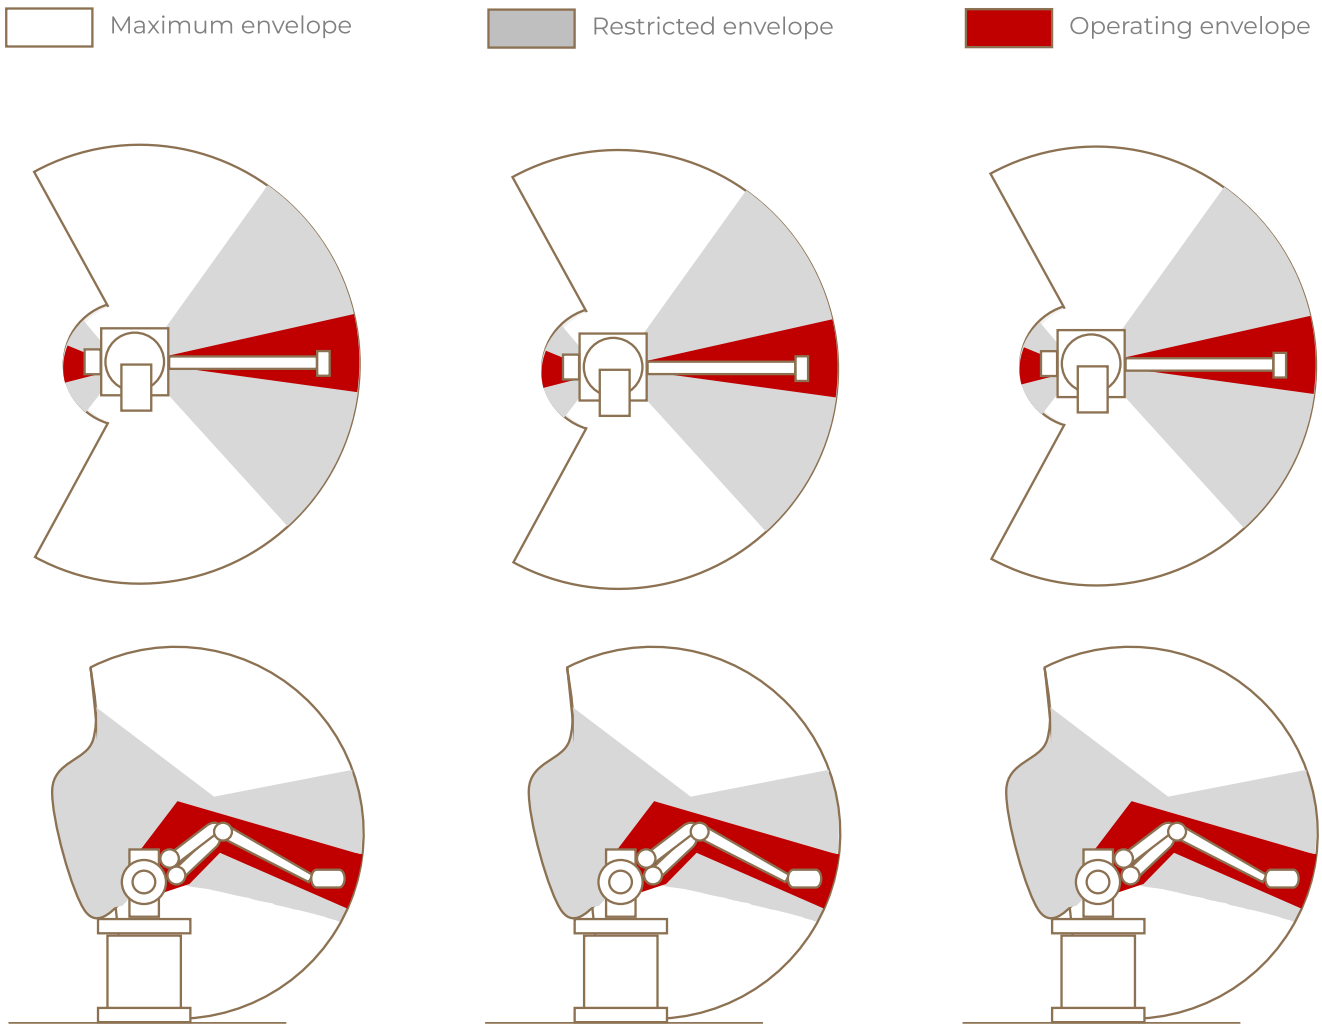

[1]

As a result, the representation of the moral situation is crucial for understanding how to allocate moral responsibility with the aim of mitigating ethical risks. The choice of a level of informational abstraction becomes particularly important with the deployment of artificial intelligence systems that have an autonomous agency capacity. In this case, moral allocation is more difficult pragmatically, but the process must be clear conceptually in order to use resources effectively.

1 https://www.distrelec.de/en/industrial-robots/cms/knowhow-industrial-robots

MORAL AGENCY

In general terms, agency is the capacity that an entity has to act[1] . An agent has agency. From an ethical point of view, the causal relationship between the capacity to act or agency and the action which ensues it simplifies the allocation of responsibility and the prescription of actions. If a human being takes an action, they have a moral capacity to act, and therefore moral responsibility falls on them. Thus a human being can be a moral agent (as the source of a moral action), and at the same time a moral patient (as the recipient of a morally loaded action). However, there are moral patients who never become moral agents. “Many entities, most notably animals, seem to qualify as moral patients, even if they are in principle excluded from playing the rôle of moral agents” (Floridi Sanders, On the Morality of Artificial Agents, p 350). Much has been done to redress this gap. “In particular, the concept of ‘moral agent’ has been stretched to include both natural and legal persons” (ibid.).
In this way, a company has a moral responsibility for the actions it takes, and this responsibility is embodied in the CEO or the board. Similarly, if a dog bites a person, the moral responsibility falls on the animal’s owner. However, according to this logic, “an entity is considered a moral agent only if (i) it is an individual agent and (ii) it is human-based, in the sense that it is either human or at least reducible to an identifiable aggregation of human beings, who remain the only morally responsible sources of action, like ghosts in the legal machine”. It follows that “insisting on the necessarily human-based nature of the agent means undermining the possibility of understanding another major transformation in the ethical field, the appearance of artificial agents (AAs) that are sufficiently informed, ‘smart’, autonomous and able to perform morally relevant actions independently of the humans who created them, causing ‘artificial good’ and ‘artificial evil’ (Gipps, 1995). Both constraints can be eliminated by fully revising the concept of ‘moral agent’” (ibid.).

The spread of artificial intelligence systems has prompted ethics to consider artificial agents as “legitimate sources of im/moral actions”. Thus **AI systems “though not intelligent and fully responsible, can be fully accountable sources of moral action”**.

An AI system is not a monolithic system, but rather a complex system in which minor agents coordinate to activate information processes (or actions). In this case, the final consequences may be more important than the sum of the actions of the individual agents. Thus “a multi-agent system (from a whole society to just a group of agents, some of which may not be human, i.e., a group of bots interacting online) may be correctly interpreted as being equivalent to a multi-layered neural network” in which even if the system produces a set of neutral actions, the result may become morally charged. In order to mitigate the risks of multi-agent systems, as in the case of AI systems, one should draw on a distributed morality approach that allows “the allocation, by default and overridable, of full moral responsibility (faultless responsibility) to all the nodes/agents in the network causally relevant for [...] independently of intentionality” causing morally loaded consequences (Floridi, Faultless responsibility).

1 https://plato.stanford.edu/entries/agency/

DISTRIBUTED MORAL ACTIONS

Upon acknowledging that a moral harm may be the unintentional result of a set of actions activated by an artificial multiagent system designed with good intentions, it may be agreed that a process to mitigate such risks is needed in any organisation. The proposed mechanism to mitigate these risks uses three concepts: "back propagation from network theory, strict liability from jurisprudence and common knowledge from epistemic logic" (Floridi, Faultless responsibility).

If we assume that the agents that compose a multi-agent system "are autonomous (in the minimal sense that they are in charge and regulate their own actions, at least to some significant extent), can interact with each other and their environments and can learn from their interactions (can change the rules according to which they behave, again, at least to some significant extent), then we can compare [... *as an analogy*] a multi-agent system to a multi-layered neural network" (ibid.).

When inputs enter the system they propagate, interacting with each other and performing morally neutral actions. But on arrival, the set of actions may cause good or bad consequences: as a result, actions become are morally-loaded.

To mitigate risks and correct consequences we must make use of a process that allows us to

a) **identify system of actions that caused morally loaded consequences (either commendable or punishable)**
b) **identify how these actions have propagated through the network (forward propagation);**
c) **allocate a maximum responsibility to each agent in the network (back propagation);**
d) **make corrections to the nodes to improve the output (of course "Some nodes may share different degrees of responsibility, including none at all, if an agent is able to show no involvement in the interactions" leading to the final consequences).**
e) **repeat steps (a) and (d) until the result is axiologically satisfactory.**

Forward propagation and Back propagation

'Forward propagation' illustrates how actions propagate through the network to consequences, while 'back propagation' is the reverse process of improving the process and mitigating ethical risks. "In forward propagation, agents in the network collectively produce a distributed action that is morally charged, activating themselves and interacting with other agents according to some specific inputs and thresholds, in ways that are assumed to be morally neutral. In such a distributed context, it no longer matters which agent does what or why"(ibid.). What is relevant is how the **allocation of moral responsibility** can be made. Even if thanks to the concept of strict liability (see below) all agents in the system are considered to be "responsible by default", it will be necessary to set up an "overridability clause". As previously noticed, "some nodes may share different degrees of responsibility, including none, if an agent is able to show no involvement in the interactions leading to the consequences"(ibid.).

In any case, the key point remains the **back propagation** process. "All that matters is that the change in the system caused by the distributed moral actions is good or bad, and if it is bad, one can try to correct or reduce it by treating the whole network as responsible for it, and then propagate back the responsibility to all its nodes/agents to improve the outcome. The cycle ends when the output is satisfactory, according to

the chosen axiological analysis" (ibid.), i.e. when the results are aligned with the moral values of the organisation. In order to distribute moral responsibility, one has to borrow the legal concept of 'strict liability' or 'faultless responsibility'.

Strict liability or faultless responsibility

In tort law, one can be held liable for one's acts or omissions, regardless of culpability. In this case, if a dog bites a bystander, the owner of the animal is liable for failing to take preventive measures.
"Under strict liability, there is no requirement to prove fault, negligence or intention. Interestingly, strict liability is most commonly associated with damage caused by animals and defectively manufactured products. [...] This establishes how far a corporation, as a legal person, can be liable for the acts and omissions of the natural persons it employs" (Ibid.).

Common knowledge

If all agents know that they will all be responsible for negative consequences, these are more likely not to occur. In practice, making every agent in the process aware of moral responsibility, perhaps through public announcements or direct training of agents, makes it possible to increase social pressure and prevent moral consequences from arising.

THE DATA ETHICIST

The role of the Data Ethicist is crucial for any organisation. Particularly for those employing AI systems that may have an impact on a significant number of individuals, communities, nation-states, society[1] or the information entities that constitute it. A Data Ethicist can either become the Chief Data Ethics Officer of an organisation or, even better, be represented by a multidisciplinary committee with the appropriate expertise. Some of the tasks, but the list is not exhaustive, of a Data Ethicist are listed below.

Identifying the moral situation by enveloping the range of action of an agent.

In order to be able to make moral decisions, it is absolutely crucial to understand what the action space of a moral agent (human agent and artificial agents) is. Identifying this space serves to understand which and how many resolutions are needed for the subsequent operations: the allocation of responsibility and ethical foresight analyses.

Lightening the burden of ethical decisions on data workers

Training people to increase their awareness and competence in ethical choices is crucial for any company. But it is thought that it is enough to train developers in data ethics for digital artefacts to be ethical. In reality, burdening developers (and only them) with the ethical choices of the company is inherently unethical. They should be instructed to know when certain choices must pass into the hands of trained professionals. A Data Ethicist is needed to support developer decisions and to argue and analyse the alignment of choices with the organisation's values.

1 SSIA's 4 Sphere's of Influence

Aligning ethical values with choices through a documented axiological analysis.

Every choice has an ethical impact. The role of the Data Ethicist is to assess what the moral threshold is and to understand what consequences are “morally negligible or morally relevant”.

Shaping the digital governance of the organisation through moral evaluation.

Ethics (and in particular soft ethics) provides tools and methods for assessing “what ought and ought not to be done over and above the existing regulation, not against it, or despite its scope, or to change it, or to by-pass it (i.e., in terms of self-regulation)”. In other words, soft ethics is post-compliance ethics that goes beyond sustainable choices, to direct the organisation towards preferable choices; “in this case, ‘ought implies may’” (Floridi, 2018).

Allocating the responsibility at the right level of abstraction and granularity.

A multiagent system is often a socio-technological system involving not only the artificial agents that make up the system put into production, but the entire ethical pipeline (from the choice of developers to the prediction quality of the models). Choosing the right level of abstraction, and even more so the right level of granularity, makes it possible to allocate responsibilities correctly, make the right moral evaluations, and implement effective ethical risk mitigation mechanisms. Moreover, in allocating responsibility, the penalty is often remembered, but the reward is almost always forgotten. Allocation of responsibility improves the ethical environment by publicly recognising positive actions.

Making ethical foresight analysis.

A Data Ethicist must 1) take into account “reasonably foreseeable misuse” (i.e. in accordance with legal jurisprudence and the EU AI Act, Article 3-13); 2) have the ability to foresee possible ethical risks on emerging technologies to avoid possible contamination with the technologies of their organisation, both those already in production and those in the design phase.

Avoiding ethical deviations.

Many people think that ethics is a burden on business. This is due to a prejudice generated by an approach that is more regulatory and legislative than managerial. In reality, ethics shapes every decision in an organisation at every level. The problem is knowing how to use ethics to maximise the quality of decision-making (i.e. in the decision-making phase), and avoiding making it a normative element with which the company and decision-makers must be compliant. However, this distrust of ethics can quickly lead the organisation to the misuse of ethics as well as to falling into the most common digital ethics deviations: ethics shopping, ethics bluewashing, ethics lobbying, ethics dumping, ethics shirking. The task of a Data Ethicist is to respect a public code of conduct and prevent the organisation from deviating from its ethical path.

GENERAL PRINCIPLES

As previously noticed, according to Floridi and Cowls (2019) and the Unified Framework of Ethical Principles for AI (Floridi et al., 2018), there are five principles that can be used to summarise the social benefit of AI:

1. Beneficence: Promoting Well-Being, Preserving Dignity, and Sustaining the Planet
2. Non-Maleficence: Privacy, Security and 'Capability Caution'
3. Autonomy: The Power to Decide (to Decide)
4. Justice: Promoting Prosperity, Preserving Solidarity, Avoiding Unfairness
5. Explicability: Enabling the Other Principles through Intelligibility and Accountability

To these AI principles we should add a few primary principles without which ('sine qua non' in Latin) an ethical approach to AI could not be initiated, that is, data security and compliance with the law, an epistemological principle 'sine qua non' the process would be biased (diversity), and an eschatological principle, which encapsulates the meaning of the data ethicist's work (semantic capital).

Data security

We live in information societies or hyperhistorical societies in which the informational component is fundamental. The risks are therefore not only physical but the security risks can affect the mere digital existence of any data entities or data subject: in fact, "Only a society that lives hyperhistorically can be vitally threatened informationally, by a cyber attack. ***Only those who live by the digit may die by the digit.***" (Floridi, 2012a, p. 130). An environment suitable for ethical reasoning cannot be developed if data security is not guaranteed first. Ethical reasoning cannot flourish in an unsafe environment. The responsibility for the cybersecurity of data and processes is not directly related to the Code of Data Ethics Code, however it should be regarded as a principle without which all other ethical activities cannot be pursued.

Data security is a *sine qua non* requirement.

Compliance with the local law, regional regulations, and international standards

If soft law can be considered a post-compliance ethics, then the compliance with the law or compliance with the relevant legal framework must be regarded as a principle without which all other action (distributed moral actions) can be illegal.
This Code of Data Ethics refers to the relevant legal framework (i.e., the EU GDPR) and the Chief Data Officer (CDO), the Data Protection Officer (DPO) or the Data Control Committee (DCC) reside in the jurisdiction of the relevant legal framework.

Compliance with the law is a sine qua non requirement.

Diversity

From an informational point of view, Diversity is the constructive foundation to avoid reinforcing biases due to uniformity of view, thus reducing the risk of AI, algorithmic or autonomous systems to produce negative consequences. The Board of Directors shall have a diversity policy that is implemented for the organisational units responsible for the data processing, however this code refers to the Code of Ethics for diversity policy details.

Semantic continuity

The principle of semantic continuity requires that all actions taken by data ethicists aim at preserving the internal (organisation) and external (society) semantic capital (Floridi, 2018). To use a metaphor, having data that has lost its meaning is like having all the keys to all the doors of a city, without knowing which door each key corresponds to. This is an ideal normative principle, because we know that meaning naturally wears out over time. However, all actions must be taken for the decay of semantic capital to be slowed down.

All the references

Bibliography

American Psychological Association. Ethical Principles of Psychologists and Code of Conduct. Including 2010 and 2016 Amendments. Washington, DC: American Psychological Association 2020 https://www.apa.org/ethics/code

Couldry, Nick and Mejias, Ulises A. (2019) "5. Data and the Threat to Human Autonomy". The Costs of Connection: How Data Is Colonizing Human Life and Appropriating It for Capitalism, Redwood City: Stanford University Press, pp. 153-184. https://doi.org/10.1515/9781503609754-007

Criado Perez, Caroline (2019). Invisible Women: Data Bias in a World Designed for Men. New York: Abrams Press. ISBN 978-1-683-35314-0. OCLC 1111651744.

Felzmann H, Villaronga EF, Lutz C, Tamò-Larrieux A. (2019) Transparency you can trust: Transparency requirements for artificial intelligence between legal norms and contextual concerns. Big Data & Society. January. doi:10.1177/2053951719860542

Floridi, L. (2012a). Hyperhistory and the Philosophy of Information Policies. Philosophy & Technology, 25, 129–131.

Floridi, L. (2013a). Distributed Morality in an Information Society. Science and Engineering Ethics, 19(3), 727–743. https://doi.org/10.1007/s11948-012-9413-4

Floridi, L. (2013b). The Ethics of Information. Oxford University Press.

Floridi, L. (2016a). Faultless responsibility: On the nature and allocation of moral responsibility for distributed moral actions. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 374(2083). https://doi.org/10.1098/rsta.2016.0112

Floridi, L. (2016b) On Human Dignity as a Foundation for the Right to Privacy. Philos. Technol. 29, 307–312. https://doi.org/10.1007/s13347-016-0220-8

Floridi, L. (2018) Semantic Capital: Its Nature, Value, and Curation. Philos. Technol. 31, 481–497. https://doi.org/10.1007/s13347-018-0335-1

Floridi, L., (2018), Soft Ethics: Its Application to the General Data Protection Regulation and Its Dual Advantage. Philos. Technol. 31, 163–167. https://doi.org/10.1007/s13347-018-0315-5

Floridi, L., & Cowls, J. (2019). A Unified Framework of Five Principles for AI in Society. Harvard Data Science Review, (1), 1–15. https://doi.org/10.1162/99608f92.8cd550d1

Floridi, L., and Taddeo, M. (2016). "What is Data Ethics?" Philosophical Transactions of the Royal Society, no. Dec 28, pp. 1-5.

Floridi, L., Cowls, J., Beltrametti, M. et al. (2018), AI4People—An Ethical Framework for a Good AI Society: Opportunities, Risks, Principles, and Recommendations. Minds & Machines 28, 689–707. https://doi.org/10.1007/s11023-018-9482-5

Floridi, L., Sanders, J. (2004). On the Morality of Artificial Agents. Minds and Machines 14, 349–379 https://doi.org/10.1023/B:MIND.0000035461.63578.9d

Levine, J. (1983). "Materialism and qualia: the explanatory gap". Pacific Philosophical Quarterly, 64: 354-361.

Loi, M., Christen, M. (2020) Two Concepts of Group Privacy. Philos. Technol. 33, 207–224. https://doi.org/10.1007/s13347-019-00351-0

Light, R., Panai, E. The Self-Synchronisation of AI Ethical Principles. DISO 1, 24 (2022). https://doi.org/10.1007/s44206-022-00023-1

Mullen, P. R., Morris, C., & Lord, M. (2017). The Experience of Ethical Dilemmas, Burnout, and Stress Among Practicing Counselors. Counseling and Values, 62(1), 37–56. https://doi.org/10.1002/cvj.12048

Panai, E., Light, R. (forthcoming) Raising the Ethical Voice of the Periphery in the Construction of AI Ethics Frameworks

Panai, E. (forthcoming) The Latent Space of Data Ethics

Panai, E. (2018). Skip! The Art of Avoiding Projects. Middletwon DE, USA: Createspace.

Popper, K., (1963), Conjectures and Refutations: The Growth of Scientific Knowledge, New York and Evanston: Harper and Row.

Ripley, B., (1996), Pattern Recognition and Neural Networks, Cambridge University Press

Robinson, W. S., (1982), “Why I Am a Dualist”, in E. D. Klemke, A. D. Kline & R. Hollinger (eds.), Philosophy: The Basic Issues, New York: St. Martin’s Press.

Skinner, B. F. (1948). “Superstition” in the pigeon. Journal of Experimental Psychology, 38, 168–172.

Tsamados, A., Aggarwal, N., Cowls, J. et al. The ethics of algorithms: key problems and solutions. AI & Soc 37, 215–230 (2022). https://doi.org/10.1007/s00146-021-01154-8

van de Poel, I., Fahlquist, J. N., Doorn, N., Zwart, S., & Royakkers, L. (2012). The Problem of Many Hands: Climate Change as an Example. Science and Engineering Ethics, 18(1), 49–67. https://doi.org/10.1007/s11948-011-9276-0

PRIVATE PART

(FORTHCOMING)

Private Part

Table of Contents

- Special categories
    - Children's Data
    - Protected categories
    - Protected variables
- Phase 1: recording and generation
    - Recording
    - Generation
- Phase 2: processing
    - Data cleaning
    - Data transformation
    - Data integration
    - Data reduction
- Phase 3: use, share and dissemination
    - Data as internal resource
    - Data as an external resource
    - De-identification & Re-identification
- Risks
    - Biased Dataset
    - Common errors
    - Discrimination & Unrepresentativeness
    - Limitation to human autonomy (data)
    - Misuse of personal data
    - Opacity and distrust
    - Re-identification
    - Re-use
    - Semantic disruption
    - Special categories discrimination
    - Toxic Combination
    - Unintended or Negligent Duplication
- Strategies for risk mitigation
    - Data Curation
        - Data Quality
        - Data Purge
        - Bias mitigation.
    - Data Sufficiency
        - Data minimisation
        - Duplicata abhorre!
        - Data Duplication Control
        - Referential Databases Design
    - Group privacy
        - WHVSV privacy
        - Inferential privacy
    - Limitation to personal data
    - Individual Autonomy
        - Consent
        - Right of access
        - Right to withdraw,
        - Right to correct,
        - Right to object
        - Transportability
    - Information Privacy
        - Anonymization (Long life privacy)
        - Pseudonymization (personal privacy)

- Semantic capital
    - Semantic consistency
- Representativeness
    - Diverse inputs
- Transparency
    - Document and maintaining registers

III. Ethics of Algorithms

- Definitions
    - Artificial Intelligence
    - Machine Learning (ML)
    - Deep Learning (DL)
- map of algorithms
- Training & Validation Data VS Production Data
    - Training Phase
    - Validation dataset
    - Test dataset
    - Production phase
    - Original or Real data
    - Synthetic data
- Risks
    - Disproportionateness
    - Limitation to human autonomy (AI) or Uncontrolled behaviour
    - Neglect of data subject
    - Overfitting
    - Vagueness or Epistemic Opacity
    - Unnecessariness
    - Unrepresentativeness
    - Untrust
- Strategies for risk mitigation
    - Care
        - Human-in-the-loop
        - Human-on-the-loop
        - Human-after-the-loop
        - Ethics Risk Analysis
        - Explainability Plus
    - Interpretability
    - Explainability
        - KPI Design
        - Register curation
    - Inclusion
        - Akaike Approach
        - Preventing Overfitting
    - Individual Autonomy
        - Start/Stop
    - Proportionality
        - Parsimony
    - Representativeness
        - Diverse Inputs
        - Multi-stakeholder Feedback
        - Bias remediation
        - Bias mitigation
    - Trustworthiness
        - Validity

        - Reliability
        - Precision
        - Accuracy
    - Utility
        - Undesired consequences & Missed opportunities
        - Necessity Assessment
        - Sufficient Condition
- Process
- Essential steps
- Best Practices
    - Necessity Assessment
    - Ethics Risk Analysis
    - Benchmarking - Baselining norms for bias remediation
    - Reliability
    - Diverse inputs and multi-stakeholder feedback
    - Application of soft law
    - Validity and KPI design