PARTICIPATORY APPROACHES TO A NEW ETHICAL AND LEGAL FRAMEWORK FOR ICT

Guidelines on Data Protection Ethical and Legal Issues in ICT Research and Innovation.

AI.

Brief Summary

Iñigo de Miguel Beriain (UPV/EHU),

This document is an abbreviated version of the part of the Panelfit Guidelines related to Artificial Intelligence (AI) Anyone using it should think of it as a basic introduction, which is in no way intended to provide sufficient knowledge of the material. Its main purpose is to prevent the reader from making serious mistakes. For better information, it is advisable to consult the full version of our Guidelines, at: https://guidelines.panelfit.eu/ai/general-exposition/

If you wish step-by-step guidance, please consult: https://guidelines.panelfit.eu/ai/step-by-step/

Some valuable cases studies can be found in: https://guidelines.panelfit.eu/ai/case-study/

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 788039. *This document reflects only the author's view and the Agency is not responsible for any use that may be made of the information it contains.*

DISCLAIMER

This document is an abbreviated version of the part of the Panelfit Guidelines related to Social Networks. Anyone using it should think of it as a basic introduction, which is in no way intended to provide sufficient knowledge of the material. Its main purpose is to prevent the reader from making serious mistakes. For better information, it is advisable to consult the full version of our Guidelines, written by .Jose Antonio Castillo Parrilla and Iñigo de Miguel Beriain (UPV/EHU).

Furthermore, one must always keep in mind that the information provided in our Guidelines does not constitute legal advice; instead, all information, content, and materials provided are for general informational purposes only. The Guidelines provide general advice around EU data protection law under the GDPR. Accordingly, the reader should be aware that the situation relevant in their specific processing context, as well as in their specific jurisdiction, may deviate from the guidance provided. Indeed, information in our Guidelines may not constitute the most up-to-date legal or other information. The legal situation in relation to data processing in the EU changes regularly. New laws and new interpretations of existing laws relevant to the topics covered by the Guidelines appear frequently and changes may not be reflected in the Guidelines.

In this regard, we would highlight, without any intention of being comprehensive, at the time of writing, the significance of the following draft EU laws to the topics covered in these Guidelines: the ePrivacy Regulation, the AI Act, the Data Governance Act, the Digital Services Act, the Digital Markets Act and the Data Act.

Where relevant at the time of writing, authors may have attempted to highlight provisions of draft laws in relation to the topics covered in these Guidelines. The reader should be aware that drafts may change and that such references may not remain valid over time. Equally, authors' choices to consider certain provisions from certain draft laws should not be taken as indicative of effort to be comprehensive in addressing all relevant provisions from all draft laws.

Readers of the Guidelines should contact their DPOs and DPAs to obtain advice with respect to any particular legal matter. No reader, user, or browser of the Guidelines should act or refrain from acting on the basis of information provided without first seeking legal advice from counsel in the relevant jurisdiction. Only DPOs and DPAs can provide assurances that the information contained herein – and your interpretation of it – is applicable or appropriate to your particular situation. Use of, and access to, the Guidelines do not create any relationship between the reader, user, or browser and the authors, reviewers, validators, or commentors, of the Guidelines.

The views expressed in, or through, our Guidelines are those of the individual authors writing in their individual capacities only – not those of EU Commission, of course. All reference to reviews, validations, or provision of comments or suggestions, refer to the personal opinions of individuals acting in their personal capacities – and do not refer to the opinions of the organisations these individuals represent or to acts of these individuals in their official capacities.

The Guidelines contain links to other third-party websites. Such links are only for the convenience of the reader, user or browser; the authors and the reviewers/validators do not recommend or endorse the contents of the third-party sites.

All liability with respect to actions taken or not taken based on the contents of the PANELFIT Guidelines are hereby expressly disclaimed.

1 Human agency and oversight

The **GDPR does not prevent any form of profiling and/or automated decision-making**. However, it provides individuals with a qualified right to be informed about it, and a right not to be subject to a decision based on purely automated decision-making, including profiling. Their right to information must be satisfied through application of the lawfulness, fairness and transparency principle. This means that, **as a minimum,** controllers have to inform the data subject that they are engaging in this type of activity, provide meaningful information about the logic involved and the significance and envisaged consequences of the profiling for the data subject.

These are some essential tips if you are considering profiling or automated decision making:

- Carry out a DPIA to consider and address the risks when you start any new automated decision-making or profiling
- Make sure that you have a legal basis to carry out profiling and/or automated decision-making, and document this in their data protection policy.
- Inform the data subjects about the profiling and automated decision-making they carry out, what information they use to create the profiles, and where they get this information from.
- Make data subjects aware of their rights. Be particularly cautious if you have obtained their data indirectly.
- Introduce additional checks if profiling/automated decision-making systems involve any vulnerable groups (including children).
- Use the minimum amount of data needed and document it through a written justification on the proportionality of all personal data used. If possible, use anonymized data.
- Create a clear retention policy for the profiles that you create and inform data subjects about it
- Do not forget that you shall guarantee the right to explainability of algorithmic decisions
- You have ensured that qualified human supervision is in place from the design phase onwards, in particular on the interpretation requirements and the effective design of the interface, and the supervisors are trained.
- You have introduced measures aimed at mitigating bias

2 Technical robustness and safety

You must concentrate in four main issues[1]: Resilience to attack and security; Fallback plan and general safety; Accuracy; Reliability and reproducibility

1. Resilience to attack and security
   Article 32 of the GDPR explicitly requires the implementation of appropriate technical and organizational measures to ensure data security. **The required security measures depend on the likely impact of an AI system malfunction.** It is up to the system designer to *predict possible unlawful processing of personal data and implement security measures that would prevent or minimize it.* . This could be through measures such as restricting the usable data sources, or prohibiting certain usage patterns though licensing terms. Data protection legal framework may complement such restrictions, but is by no means a replacement for them. Thus.
   - Assess potential forms of attacks to which the AI system could be vulnerable.
   - Consider different types and natures of vulnerabilities, such as data pollution, physi **The required security measures depend on the likely impact of an AI system malfunction**cal infrastructure and cyber-attacks.
   - Put measures or systems in place to ensure the integrity and resilience of the AI system against potential attacks.
   - Verify how the system behaves in unexpected situations and environments.
   - Consider to what degree the system could be dual-use. If so, take suitable preventative measures against this (e.g. not publishing the research or deploying the system).

2. Fallback plan and general safety
   The GDPR requires controllers to implement suitable fallback plans protecting data subjects, including the right to contest an AI decision and to obtain a human intervention that considers the data subjects' point of view. Such safeguards should be considered during the systems design. Thus:
   - Introduce measures aimed at identifying potential safety risks of (other) foreseeable uses of the technology, including misuse.
   - Create a plan to mitigate or manage these risks
   - Perform a DPIA prior to processing when there is a high risk to the rights and freedoms of a natural person

[1] These tips have been adapted from the checklists elaborated by the High-Level Expert Group on Artificial Intelligence (2019) Ethics guidelines for trustworthy AI. European Commission, Brussels. Available at: https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai (accessed 20 May 2020).

- Make sure that the system has a sufficient fallback plan if it encounters adversarial attacks or other unexpected situations (e.g. technical switching procedures or asking for a human operator before proceeding)
- Consider an insurance policy to deal with potential damage from the AI system
- Assess whether there is a probable chance that the AI system may cause damage or harm to users or third parties, its likelihood, potential damage, impacted audience and severity.
- Consider the liability and consumer protection rules and the impact or safety risk to the environment or to animals.
- Perform a risk analysis that includes whether security or network problems (e.g. cybersecurity hazards) could pose safety risks or damage due to unintentional behaviour of the AI system

3. Accuracy

Many AI systems require accurate and reliable training data to achieve the best results. Thus:

- Assess what level and definition of accuracy would be required in the context of the AI system and use case
- Implement measures aimed at assessing how accuracy is measured and assured
- Verify what harms would be caused if the AI system makes inaccurate predictions and introduce measures to mitigate them
- Performs periodic checks aimed at uncovering and mitigating impact among different societal groups.
- Do not only concentrate in accuracy, but also in fairness and representativity of data and results

4. Reliability and reproducibility

AI systems it may evolve over time and slowly drift away from the designers' original intentions. It is therefore important to document clearly the initial assumptions and conditions under which the AI system was intended to be used. If an AI system is publicly available, the documentation of the system's reliability should be as well. In addition, reproducibility is also an important prerequisite for trust. If a result cannot be reproduced, its explainability - and therefore trust in the AI system - may suffer. Thus:

- Implement a strategy to monitor and test if the AI system is meeting its goals, purposes and intended applications.
- Test whether specific contexts or particular conditions need to be taken into account to ensure reproducibility.
- Put in place verification methods to measure and ensure different aspects of the system's reliability and reproducibility.
- Introduce processes to describe when an AI system fails in certain settings.

- Document and operationalize these processes for the testing and verification of the reliability of AI systems.
- Establish mechanisms of communication to assure (end-)users of the system's reliability.

3 Privacy and data governance

3.1 Purpose limitation

Controllers must limit the use of personal data to the original purpose(s), or those purposes that are compatible with it. The **re-use of data** in the development of an AI tool entails deeply challenging issues in terms of purpose limitation. AI systems process personal data in various stages and for a variety of purposes. As a result, a controller may fail to distinguish each distinct processing operation and process data for purposes others than those for which they were initially collected. Controllers should be particularly concerned about these situations since they could lead to a failure to comply with the data protection principle of lawfulness. If the AI developers are planning to use a dataset at different stages (e.g. training, validation or deployment), they should consider these steps as having distinct and separate purposes. Thus, you shall:

- Identify your purpose or purposes for processing and make sure that these are limited in time and scope
- Document those purposes.
- Include details of their purposes in the privacy information for individuals, ensuring that the data subject is adequately informed, according to art. 12-14 GDPR.
- Include details of your purposes in the privacy information for individuals.
- Regularly review your processing and, where necessary, update your documentation and privacy information for individuals.
- If you to intend to use personal data for a new purpose other than a legal obligation or function set out in the law, check that this is compatible with your original purpose or a new legal basis allowing such processing[2].

3.2 Lawfulness

Controllers shall ensure that they have a **legal basis for processing personal data**. **If this is not the case,** the processing **must not be carried out.** . In the case of AI, the legal bases that are usually invoked to justifying processing are consent, legitimate interest or public interest.

[2] ICO (no date) Principle (b): purpose limitation. Information Commissioner's Office, Wilmslow. Available at: https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/principles/purpose-limitation/ (accessed 17 May 2020).

3.2.1 Consent

Data processing is often based on consent provided by the data subjects. However, consent does not fit well with the essential nature of most AI developments. If processing involves the use of a complex AI tool that may have further uses of the data (e.g. profiling and automated decision-making might happen inadvertently, data are likely to be inferred during processing, such inferred data can be used for various purposes, etc.), it is difficult to see how a single consent could justify all such processing. To this end, controllers must take good care of the guidelines on consent provided by the EDPB. Broad consent can serve well in the context of scientific research. However, it is not a kind of blanket or equivocal to open consent. It is an **exceptional tool that can only be acceptable if several conditions apply** Thus:

- Check that consent is the most appropriate legal basis for processing.
- Request the consent of the interested parties in a free, specific, informed and unequivocal manner.
- Make sure that you strictly follow the best practices guidelines related to consent when children or people unable to provide consent are involved
- Make sure that broad consent is used for (and only for) scientific research, and only when it is difficult or improbable to foresee how this data will be processed in the future.
- Assess that broad consent used for processing of special categories of data is compatible with national regulations.
- Where broad consent is used, the data subjects are given the opportunity to withdraw their consent and to choose whether or not to participate in certain research and parts of it.
- Build a direct relationship with the subject who provides the data to be used for training, validation and deployment of the IA model.
- Make sure that there is no power imbalance between controllers and data subjects.
- Ask people to positively opt in.
- Do not use pre-ticked boxes or any other type of default consent.
- Use clear, plain language that is easy to understand.
- Specify what kind of data you collect, why you want the data, what you are going to do with it and for how long.
- Give separate distinct (‘granular’) options to consent separately to different purposes and types of processing.
- Tell individuals they can withdraw their consent and how to do so.
- Ensure that individuals can refuse to consent without detriment.
- Avoid making consent a precondition of a service.

3.2.2 Legitimate interest

The use of legitimate interest as a legal ground for processing for AI development is applicable, provided that the result of the balancing test justifies it. This **should be adequately documented in the records of processing.**

- Check that legitimate interest is the most appropriate basis.
- Be aware that you are responsible for the protection of individuals' interests.
- Keep a record of the decisions made and the reasoning behind them, to ensure that you can justify your decisions.
- Identify the relevant legitimate interests.
- Check that the processing is necessary and there is no less intrusive way to achieve the same result.
- Perform a balancing test and are confident that the individual's interests do not override those legitimate interests and document it.
- Only use individuals' data in ways they would reasonably expect, unless you have a very good reason to do otherwise. This reason should be documented and available to third parties
- Make sure that you are not using people's data in ways they would find intrusive, or which could cause them harm, unless you have a very good reason. These reasons should be documented and available to third parties
- If you process children's data, they take extra care to make sure they protect the children's interests.
- Considered whether you also need to conduct a DPIA.
- Consider introducing safeguards to reduce any possible negative impact detected by the DPIA, where possible.
- Consider whether your can offer an opt out.

3.2.3 Public interest

Public interest is an essential legal basis for data processing and scientific research is considered to be "in public interest". The Regulation envisages a special and favorable regime to the processing operations for this purpose. This includes a flexible regime for long-term data storage and a presumption of compatibility for secondary or further purposes. In addition, limitations, exceptions, or derogations are provided, *inter alia*, to the rights to information, access, rectification, restriction of processing, object and, concerning the archiving purposes in the public interest, to the right of notification and portability. In order to strike the right balance with data subject's rights and interests, the Regulation requires the adoption of appropriate safeguards in accordance with Article 89 and, in certain situations, also further development by Union or Member State law. Thus:

- You ensure that your research can be considered scientific research
- You implement the safeguards foreseen by the Regulation or the rules by your Member State
- You pay special attention to derogations according to the GDPR.

Please, dedicate some time to consult the part of our Guidelines devoted to this issue, at: https://guidelines.panelfit.eu/the-gdpr/main-concepts/data-protection-and-scientific-research/

3.3 Data minimization

The data minimization principle requires AI developers to opt for those tools whose development involves minimal use of personal data compared to the available alternatives. **Controllers should avoid using personal data if it is not necessary**; that is, if the objective that the controller is aimed at can be obtained without processing personal data. If anonymization is not possible, controllers should at least try to work with pseudonymized data. They must be able to demonstrate that the processing is **necessary for the objective being pursued** and is **less intrusive than other options** for achieving the same goal; not that it is a necessary part of their chosen methods. Some additional measures related to the minimization principle include:

- limit the extension of the data categories (e.g. names, physical and addresses, fields about their health, work situation, beliefs, ideology, etc.)
- limit the degree of detail or precision of the information, the granularity of the collection in time and frequency, and the age of the information used
- limit the extension in the number of interested parties of those who treat the data
- limit the accessibility of the different categories of data to the staff of the controller/manager or even the end-user (if there are data from third parties in the AI models) at all stages of the processing.

Thus, some important tips are:

- Make sure that you only use personal data if needed.
- List all data sources and provided justification for their need and proportionality in the context of this study/project and document it.
- Considered the proportionality between the amount of data and the accuracy of the AI tool and document your conclusions (and, of course, do not process the data if there is no proportionality).
- Periodically review the data you hold, and delete anything they do not need.
- the controllers have)”
- At the training stage of the AI system, delete all information not strictly necessary for such training.
- Check if personal data are processed at the distribution stage of the AI system and delete them unless there is a justified need and legitimacy to keep them for other compatible purposes.

4 Fairness with respect to data subjects’ rights

4.1 Right to information

The controller should provide the data subjects with complete information about the processing and their rights in an understandable format. Some essential tips include:

What to provide:

- You are aware that you need to inform individuals of their right to information, in addition to including it in your privacy notice.
- If the personal data were directly provided by the data subject, provide all the information enlisted in Article 13.1 GDPR;
- If the personal data were not provided by the data subject, provide all the information enlisted in Article 14.1 – 2 GDPR;If the information were already fully provided to the data subject, no need to comply with this obligation anymore.

When to provide:

- At the time the information was collected from the data subject;
- When the data are not collected from the data subject:
- Within a reasonable period after obtaining the personal data, but at the latest within one month;
- If the personal data are to be used for communication with the data subject, at the latest at the time of the first communication to that data subject;
- If a disclosure to someone else is envisaged, at the latest when the personal data are first disclosed.

How to provide:

- Concisely;
- Transparently;
- Intelligibly;
- Easily accessible;
- In a clear and plain language.

Exemptions:

- When the data subject already has all the relevant information;
- If the personal data were not provided by the data subject:
- When the provision of information is impossible or disproportionate.

4.2 Right to access

Data subjects' right to access their data must be **guaranteed in all the steps of an AI tool's life cycle. This includes observed, derived and inferred data**

Some essential tips are:

- You are aware that you need to inform individuals of their right to access, in addition to including it in your privacy notice.
- Ensure that you have provided data subjects with clear information on how to exercise their access rights

- Make sure that you are able to recognize a subject access request and they understand when the right of access applies.
- Understand that the right of access is to be applied at each stage of the life cycle of the AI solution, if it uses personal data.
- Have a policy for how to record requests you receive verbally.
- Understand when you can refuse a request and be aware of the information they need to provide to individuals when doing so.
- Understand the nature of the supplementary information you need to provide in response to a subject access request.
- Have processes in place to ensure that you respond to a subject access request without undue delay and within one month of receipt.
- Be aware of the circumstances in which you can extend the time limit to respond to a request.
- Consider that there is a particular emphasis on using clear and plain language, especially if you are disclosing information to a minor. Consider who should be the subject of the information (the child? A representative?)
- Understand what you need to consider if a request includes information about others.
- Understand how to apply the right to access in training stages.

4.3 Right to data portability

This right provides data subjects with control of the use of their data by redirecting it where it is most useful. However, it only covers data provided by the data subject. Thus, it excludes ‘inferred data’ and ‘derived data’, that is, personal data that are created by a service provider (e.g. algorithmic results). Different to observed or gathered data, **inferred data are created by the service itself, based on the observed data, not provided by the data subject. Some fundamental tips are:**

- You are aware that you need to inform individuals of their right to portability, in addition to including it in your privacy notice.
- Take into account the requirement for data portability from the earliest stages of conception and design of the AI processing. Otherwise, things will get seriously complicated if a data subject ask for this right.
- Make sure that you are able to recognize a request for data portability and understand when the right applies.
- Be aware of the circumstances that allow you refuse a request and be aware of the information you need to provide to individuals if you proceed with such refusal.
- If the portability request is made by several data subjects, make sure that all of them agree on the request
- If the information intertwines with the one from other individuals, please carry out a balancing test.
- Transmit data in structured, commonly used and machine-readable formats;

- The controllers inform users in advance when it is not technically possible to exercise the right of portability by means of a protocol.
- Transmit data in a secure way.
- Implement adequate processes to ensure that you respond to a request for data portability without undue delay and within one month of receipt. If it is going to take longer, inform the data subject about the delay and the time it will take to process the request.
- Be aware of the circumstances under which you can extend the time limit to respond to a request.

4.4 Right to restriction

Article 18 GDPR enables the data subjects to temporarily restrict a controller from processing their personal data. As provided by Recital 67 GDPR, the **methods** in which the controller can restrict personal data processing can include, for example, temporary movement of the selected data to another processing system, making the data unavailable to users or the removal of personal data on a temporary basis. Overall, the aim is to prevent data from being processed, with the exception of the storage (Article 18.2 GDPR).

Some essential tips are:

- If you receive a request to restrict data the processing from a legal entity, please indicate that the request was not lodged by an individual;
- If the individuals have not identified themselves, please ask for further information to confirm identity
- If the request does not fall within one of the scenarios laid down in Article 18.1 GDPR, please inform the data subject that the request shall be denied
- If the request cannot be fulfilled within one month, please inform why and how long will it take to process the request
- Remember that the restriction does not encompass the data storage;
- When restriction is pending, personal data can still be processed under the circumstances laid down in Article 18.2 GDPR;
- Communicate the restriction of the processing to each recipient to whom the personal data has been disclosed in compliance with Article 19 GDPR, unless this proves impossible or involves disproportionate effort.

4.5 Right to rectification

The right to correct inaccurate data is particularly important in the case of AI, since machine learning algorithms often infer data. It is particularly important to keep in mind that if the controller finds that, contrary to the views of the data subject, the data is not inaccurate with regard to the purposes of processing, the controller **does not have to rectify the data.** However, the burden of the proof is placed in the controllers' shoulders. They must provide a good reason to deny rectification, and it is hard to conclude that the damage this could bring to the AI system might serve as a convincing reason. The EDPS has criticized systems that do not include the option to have a set of individual personal data rectified without creating considerable harm to the whole system. In any case, if the controller opts to deny the data subjects' request, they must reply to the data subject with a justified reason for not rectifying the data and, if they wish to, the data subject can then refer the matter to the supervisory authority. Some essential tips are:[3] Ensure that

- You are aware that you need to inform individuals of their right to rectification, in addition to including it in your privacy notice.
- You know how to recognize a request for rectification and understand when this right applies.
- Did you receive a rectification request from a legal entity? If yes, please indicate that the request was not lodged by an individual;
- If the data subjects have not identified themselves in an adequate manner, please ask for further information to confirm identity
- You have a policy for how to record requests you receive (including verbally).
- You understand when you can refuse a request, and you are aware of the information you need to provide to individuals when asked to do so.
- Do you need a proof of inaccuracy or additional information to rectify the data? If yes, please ask for further information to the data subject. Remember not to place an unreasonable burden of proof on the data subject
- You are prepared to address the right of rectification of data subjects' data, especially those generated by the inferences and profiles made by the AI solution.
- You have processes in place to ensure that they respond to a request for rectification without undue delay and within one month of receipt.
- You are aware of the circumstances when they can extend the time limit to respond to a request.
- You have appropriate systems to rectify or complete information, or provide a supplementary statement.
- You have procedures in place to inform recipients if you rectify any data you have shared with them, unless this proves impossible or involves disproportionate effort.

4.6 Right to erasure

Data subjects have a permanent right to ask the controller for the deletion of their personal data. This might be extremely complicated in some cases, however. Indeed,

[3] These have been created on the basis of ICO (no date) Right to rectification. Information Commissioner's Office, Wilmslow. Available at: https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/individual-rights/right-to-rectification/ (accessed 28 May 2020).

one must keep in mind that sometimes it may be impossible to fulfil the legal aims of the right to erasure – also known as the right to be forgotten – in AI environments, since the obscurity of the processing might hide some personal data event to the processor. Furthermore, some data might be essential for the AI tool. Thus, controllers should always try to arrive at a simple conclusion: **the best way to avoid catastrophic damage is to prepare for a possible loss of data from the very beginning.** Some essential tips are[4]

- You are aware that you need to inform individuals of their right to erasure, in addition to including it in your privacy notice.
- You know how to recognize a request for erasure and they understand when the right applies (see article 17.1 GDPR).
- You are aware that if the request satisfy one of the exemptions provided by Article 17.3 GDPR you can inform and explain to the data subject that the request shall be denied
- You have a policy for how to record requests that you receive (even verbally).
- You understand when you can refuse a request and are aware of the information you need to provide to individuals when doing so.
- You have processes in place to ensure that you respond to a request for erasure without undue delay and within one month of receipt.
- You are aware of the circumstances under which you can extend the time limit to respond to a request.
- You understand that there is a particular emphasis on the right to erasure if the request relates to data collected from minors
- You have procedures to inform recipients if they erase any data you shared with them, unless this proves impossible or involves disproportionate effort
- You have appropriate methods to erase information in robust, accountable and permanent way, which prevents you and any other party from (re-) accessing and (re-)processing the data;

4.7 Right to object

Data subjects have the right to object to the processing of their personal data when the controller processes them on the basis of a legitimate interest, or for a task in the public interest. This does not apply to cases where the legal ground for processing was informed consent, since in those cases data subjects could simply withdraw their consent and the controller could no longer process their data. Once data subjects make their request, controllers must cease to process the data, unless they can prove they have compelling and justifiable grounds for continuing to do so, and that these grounds outweigh the data subjects' interests, rights and freedoms. Some essential tips are:

[4] ICO (no date) Right to erasure. Information Commissioner's Office, Wilmslow. Available at: https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/individual-rights/right-to-erasure/ (accessed 28 May 2020).

- You have clear information in their privacy notice about individuals' right to object, which is presented separately from other information on their rights.
- You understand when you need to inform individuals of their right to object, in addition to including it in their privacy notice.
- You know how to recognize an objection and they understand when the right applies.
- You are aware of the fact that if the request falls within one of the exceptions laid down in Article 21.2-6 GDPR, you shall inform the data subject that the request shall be denied.
- You have a policy for how to record objections you receive (even verbally).
- You understand when they can refuse an objection and are aware of the information they need to provide to individuals when doing so.
- You have processes in place to ensure that you respond to an objection without undue delay and within one month of receipt.
- You are aware of the circumstances when you can extend the time limit to respond to an objection.
- You are able to check the data subject's particular situation aim at balancing its rights with the legitimate ones of others in processing their data.

5 Transparency

Transparency is a key prerequisite for accountability in the GDPR. Its main focus is to inform data subjects upfront of the existence of the processing and its main characteristics. Other information (e.g. about the data subject) is available on request. Data subjects must also be informed of certain events, notably data breaches (in cases where the data subject is exposed to high risk). Under the GDPR, data subjects are empowered to be the main guardians of their own rights and freedoms. Evidently, transparency is a prerequisite for detecting and intervening in case of non-compliance.

Thus, transparency means that data subjects are provided with clear information about data processing. They must be informed about how and for which purposes their information (including both observed and inferred data about them) is used, no matter whether this information is collected from the data subjects themselves or by others. Data subjects should always be aware of how and why an AI-assisted decision about them was made, or where their personal data was used to train and test an AI system. Controllers must keep in mind that in such cases, transparency is even more important than when they have no direct relationship with the data subjects.

In general, transparency must be guaranteed by using a number of complementary tools. Naming a DPO, who then serves as a single point of contact for queries from data subjects, is an excellent option. Preparing adequate records of processing for the supervisory authorities, or performing DPIAs, are also highly recommended measures to promote transparency. And undertaking analyses that evaluate the effectiveness and accessibility of the information provided to the data subjects helps to ensure the efficient implementation of this principle.

In general, controllers should always provide for the development of more understandable algorithms over less understandable ones. Trade-offs between the **explainability, transparency and best performance of the system** must be appropriately balanced based on the context of use. If controllers have no choice but to use an opaque model, they should at least try to find technical solutions to the lack of interpretability.

Some important tips include[5]:

Traceability

Establish measures that can ensure traceability. This could entail documenting the following methods:

- Methods used for designing and developing the algorithmic system:
    - Rule-based AI systems: the method of programming or how the model was built;
    - Learning-based AI systems; the method of training the algorithm, including which input data was gathered and selected, and how this occurred.
- Methods used to test and validate the algorithmic system:
    - Rule-based AI systems; the scenarios or cases used in order to test and validate;
    - Learning-based model: information about the data used to test and validate.
- Outcomes of the algorithmic system:
    - The outcomes of or decisions taken by the algorithm, as well as potential other decisions that would result from different cases (for example, for other subgroups of users).

Explainability:

- Assess:
    - to what extent the decisions and hence the outcome made by the AI system can be understood?
    - to what degree the system's decision influences the organisation's decision-making processes?
    - why this particular system was deployed in this specific area?
    - what the system's business model is (for example, how does it create value for the organisation)?

[5] These tips have been adapted from the checklists elaborated by the High-Level Expert Group on Artificial Intelligence (2019) Ethics guidelines for trustworthy AI. European Commission, Brussels. Available at: https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai (accessed 20 May 2020).

- Ensure an explanation as to why the system took a certain choice resulting in a certain outcome that all users can understand?
- Design the AI system with interpretability in mind from the start?
  - Research and try to use the simplest and most interpretable model possible for the application in question
  - Assess whether you can analyse your training and testing data Can you change and update this over time
  - Assess whether you can examine interpretability after the model's training and development, or whether you have access to the internal workflow of the model

Communication:

- Communicate to (end-) users – through a disclaimer or any other means – that they are interacting with an AI system and not with another human Did you label your AI system as such
- Establish mechanisms to inform (end-)users on the reasons and criteria behind the AI system's outcomes
  - Communicate this clearly and intelligibly to the intended audience
  - Establish processes that consider users' feedback and use this to adapt the system
  - Communicate around potential or perceived risks, such as bias
  - Depending on the use case, consider communication and transparency towards other audiences, third parties or the general public
- Clarify the purpose of the AI system and who or what may benefit from the product/service.
  - Specify usage scenarios for the product and clearly communicate these to ensure that it is understandable and appropriate for the intended audience
  - Depending on the use case, think about human psychology and potential limitations, such as risk of confusion, confirmation bias or cognitive fatigue
- Clearly communicate characteristics, limitations and potential shortcomings of the AI system
  - In case of the system's development: to whoever is deploying it into a product or service
  - In case of the system's deployment: to the (end-)user or consumer

6 Fairness, diversity and non-discrimination

In the AI field, biases constitute a formidable threat against this principle, because they could lead to potential stigmatization or discrimination of isolated individuals or entire communities. Algorithms' development processes **should always include a careful monitoring of possible biases.** Internal and external reviews should pay special attention to this issue. Datasets built for validation purposes should be carefully selected to ensure an adequate incorporation of data pertaining to subjects from different sectors of society, in terms of age, race, gender, disabilities, etc. Fortunately, there are a lot of technical tools devoted to eradicating biases in AI models. The IEEE P7003TM Standard for Algorithmic Bias Considerations is particularly interesting at the moment. Some essential tips are:

- Establish a strategy or a set of procedures to **identify** bias in data selection, analysis and processing
- Implement a strategy or a set of procedures to avoid creating or reinforcing unfair bias in the AI system, both regarding the use of input data and for the algorithm design.
- Assess and acknowledge the possible limitations stemming from the composition of the used datasets and do whatever possible to fix the issues found.
- Consider the diversity and representativeness of the data used.
- Identify and protect in specific, accountable ways the vulnerable populations in the dataset by testing the tool for those specific populations and the problematic use cases.
- Use the available technical tools to improve their understanding of the data, model and performance.
- Put in place processes to test and monitor for potential biases during the development, deployment and use phases of the AI system.
- Implement a mechanism that allows others (inside and outside his/her organization) to flag issues related to bias, discrimination or poor performance of the AI system, both in the development and in the deployment stages, if possible.
- Establish clear steps and ways of communicating how and to whom such issues can be raised.
- Considered others, potentially indirectly affected by the AI system, in addition to the (end)users.
- In case of variability, establish a measurement or assessment mechanism of the potential impact of such variability on fundamental rights.
- Implement a quantitative analysis or metrics to measure and test the applied definition of fairness.

7 Societal and environmental well-being

The GDPR does not include specific provisions related to societal and environmental well-being. However, Article 5(1)(b) states that “personal data shall be collected for specific, explicit and legitimate purposes”. Through this clause, the GDPR introduces the concept of legitimacy in the data protection context. In the ‘White paper on artificial intelligence: a European approach to excellence and trust', the authors note that "[g]iven the increasing importance of AI, the environmental impact of AI systems needs to be duly considered throughout their lifecycle and across the entire supply chain, e.g. as regards resource usage for the training of algorithms and the storage of data".[6]

Further concrete recommendations for AI development that are oriented to societal and environmental well-being can be found in the 'Report from the Commission to the European Parliament, the Council and the European economic and social committee: report on the safety and liability implications of artificial intelligence, the internet of things and robotics'.[7] This kind of ethical recommendations should be carefully considered by AI developers before processing personal data, since they are clearly linked to their legitimacy.

8 Accountability

"The requirement of accountability complements the above requirements, and is closely linked to the principle of fairness. It necessitates that mechanisms be put in place to ensure responsibility and accountability for AI systems and their outcomes, both before and after their development, deployment and use."

- High-Level Expert Group on AI[8]

8.1 Accountability[9]

Accountability consists of two requirements for controllers:

- **Compliance** with the principles of the GDPR;
- **Demonstration of compliance.**

Compliance is achieved by implementing ***technical and organizational measures*** that are adequate compared to the risks to the rights and freedoms of data subjects, correspond to the state of the art of technology, and are cost-effective. Every description of the principles has provided examples of such technical and organizational measures.

[6] European Commission (2020) White Paper on artificial intelligence: a European approach to excellence and trust. European Commission, Brussels, p.3. Available at: https://ec.europa.eu/info/sites/info/files/commission-white-paper-artificial-intelligence-feb2020_en.pdf (accessed 26 May 2020).

[7] European Commission (2020) Report from the Commission to the European Parliament, the Council and the European economic and social committee: report on the safety and liability implications of artificial intelligence, the internet of things and robotics. European Commission, Brussels. Available at: https://ec.europa.eu/info/sites/info/files/report-safety-liability-artificial-intelligence-feb2020_en_1.pdf (accessed 26 May 2020).

[8] High-Level Expert Group on Artificial Intelligence (2019) Ethics guidelines for trustworthy AI, p.19. European Commission, Brussels. Available at: https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai (accessed 15 May 2020).

[9] This part of these Short Guidelines was originally written by Bud Brugger, see: https://guidelines.panelfit.eu/the-gdpr/main-principles/accountability/

For a systematic application of these measures, controllers can create ***data protection policies***. ***Approved codes of conduct***, where available, are similar but are pre-approved and usually address an entire sector. Compliance is not a state that is reached once, but **a continuous process** that spans the whole life cycle of a processing activity.

Demonstration of compliance is predominantly achieved by **documentation**. Documentation should be continuous like the process of compliance. Every implemented measure, including data-protection-relevant considerations and decisions, should be documented. The GDPR requires two formal documents as part of demonstrating compliance towards supervisory authorities: the **register of processing** and, where the risks are likely to be high, a **data protection impact assessment.** Certification can support the demonstration of compliance.

8.2 Documenting of processing[10]

The main decisions made by the data controller "have to be documented in order to comply with the requirement of data protection by design" (of Article 25 of the GDPR). Indeed, an organization who is processing personal data (including both, controllers and processors) needs to documents its activities primarily for consumption by the competent Data Protection Supervisory Authorities (DPA). This includes the ***records of processing*** that is maintained centrally by the organization across all its processing activities and **additional documentation** that pertains to an individual data processing activity.

Records of processing can be kept in written or electronic form[1]. So expect to either fill in an organization-specific form or enter your information into some (data protection) management system.

To provide an initial idea, the minimal content of the records of processing for controllers includes the following items[2]:

- The **name** and **contact** details **of the controller**, the controller's **representative** and the **data protection officer**;
- the **purposes** of the processing;
- a description of the **categories of data subjects** and of the **categories of personal data**;
- the **categories of recipients** to whom the personal data have been or will be disclosed;
- where applicable, **transfers of personal data to a third country** together with the documentation of suitable safeguards;
- where possible, the envisaged **time limits for erasure of the different categories of data**
- where possible, a general description of the **technical and organizational *security* measures**

Keep in mind that:

[10] This part of these Short Guidelines was originally written by Bud Brugger, see: https://guidelines.panelfit.eu/the-gdpr/main-tools-and-actions/documentation-of-processing-personal-data/

- Your organization may use a different set of items since on one hand, it already is in possession of some of this information (such as the first bullet), and on the other hand, it may require additional information (such as the contact of the person responsible for the single processing activity at hand).
- It is possible that the legally required record keeping is combined with the management needs of the organization, such as an internal inventory of computing and computing resources.
- Your organization may also use multiple systems, e.g. depending on whether it is acting as a controller or as a processor; or distinguishing between permanent data processing activities (such as communication systems and accounting) and temporary ones (such as those linked to temporary projects or assignments). The creation and maintenance of records across multiple systems is not prohibited under the GDPR.

Some essential tips are:

- Data protection (like security) is a process, not a final state. Continuously document that process rather than only the final characteristics of the processing activity.
- When applying data protection by design[4], the processing activity can be seen as the results of a series of many considerations and decisions. It is these considerations and decisions that should be documented.
- Deciding on a structure and format to systematically collect this information at the point of time when you conceive your processing activity.
- Where the documentation itself contains personal information (see below), make sure to protect is sufficiently and limit its further use to the purpose of demonstrating compliance with the GDPR.

8.3 Risk assessment and DPIAs

A DPIA is a process in which the data controller, before starting a data-processing procedure with high risk to the fundamental rights and freedoms of data subjects, assesses the impact of the envisaged processing operations on the protection of personal data (Article 35(1)). Determining if the data processing is of high risk is not an easy task, however. Article 35(3) lists three cases: (1) a systematic and extensive evaluation of personal aspects relating to natural persons, which is based on automated processing, including profiling, and on which decisions are based that produce legal effects concerning the natural person or similarly significantly affect the natural person; (2) processing on a large scale of special categories of data referred to in Article 9(1), or of personal data relating to criminal convictions and offences referred to in Article 10; and (3) a systematic monitoring of a publicly accessible area on a large scale.

Recital 90 of the GDPR further clarifies that the assessment of risk should be done using two parameters: the **likelihood** and **severity** of high risk, taking into account the nature, scope, context and purposes of the processing and the sources of risk. Some concrete tips:

- Be aware of the jurisdictions where data-processing activities will take place.

- Check if those jurisdictions have enacted lists indicating the processing activities that require a DPIA and check if the intended data processing activities that involve AI are covered by those provisions.
- If you are unsure of the necessity of carrying out a DPIA, consult with the DPO or, in lieu of, your legal department.
- If necessary, carry out a DPIA.
- If necessary, file a prior consultation with the appropriate supervisory authority.
- If changes are suggested, follow the advice of the supervisory authority

In order to see if a DPIA is necessary:

- Determine the jurisdictions where data-processing activities will take place.
- Check if those jurisdictions have enacted lists indicating the processing that requires a mandatory DPIA and checked if the intended data processing is covered by those provisions.
- If you are unsure of the necessity of carrying out a DPIA, you must consult with the DPO or, in lieu of, the legal department of the controller.
- If necessary, file a prior consultation with the appropriate supervisory authority.

There is no standard way to perform a DPIA. However, Article 35.7 GDPR calls for specific elements that shall always be present. These are:

- a systematic description of the envisage processing operations;
- the purposes of the processing operations;
- an assessment of the necessity of the processing operations in relation to the purposes;
- an assessment of the proportionality of the processing operations in relation to the purposes;
- an assessment of the risks to the rights and freedoms of data subjects;
- the technical and organizational measures envisaged to address the risks.

8.4 DPOs

In the case of AI development **the appointment of a DPO is (almost) certainly necessary, as many AI systems process personal data, which would make them applicable under the conditions described in Article 37(1)(a) and (b) in most situations.** If a DPO has to be appointed, for any of the reasons mentioned above, it is necessary to have their participation in the DPIAs as well as any other issue related to data protection within the entity (as prescribed by Article 39(1)(a)). This may include reviewing a potential processor, as described in the previous item. Therefore, the researchers involved in the development of the AI should consult with the DPO regarding the data-protection issues that might arise during the development of the AI. Thus, the tips are:

- Appoint a DPO if possible and whenever it constitutes a legal requirement
- If your are not required to appoint it and you do not proceed to name it on a voluntary basis, introduce the necessary organisational and security measures
- Make sure that the DPO is aware of every step taken to allow room for their intervention if deemed relevant and able to carry them out.

8.5 Processor due diligence

The accountability principle is also present when a controller chooses to require the services of a processor. Therefore, a researcher conducting AI development that needs to hire a third party for certain processing activities would need to ask two questions: (1) what type of conduct is expected to demonstrate compliance with this obligation; and (2), if some form of positive action is expected, how should controllers proceed to carry such due diligence. In general, **any question that You would ask themselves when developing the AI should be asked to the processor.** These comprise:

- Require information regarding where the data-processing activities will take place, and: (1) carry out the case law review suggested below; and (2) assess if the jurisdictions, in case of non-EU countries, are deemed as adequate by the EU Commission.
- Review case law from the national supervisory authorities where the processor operates to check for potential sanctions.
- Require proof of adherence to a code of conduct or certification.
- Require proof of relevant ISO certification.
- Require a copy of records of processing activities.
- Enquire about the development process of the AI, in particular which kind of data were used for training the AI and the data that the AI needs to operate and deliver a useful result.
- Require the processor proof of adherence to a code of conduct or certification (this is not strictly necessary but may be considered as good practise).
- Require also proof of relevant ISO certification (this is not strictly necessary but may be considered as good practise).