TABULA RASA

Trusted Biometrics under Spoofing Attacks

http://www.tabularasa-euproject.org/

Funded under the 7th FP (Seventh Framework Programme)
Theme ICT-2009.1.4
[Trustworthy Information and Communication Technologies]

D7.5: Ethical Guidelines

Due date: 30/06/2013
Submission date: 28/06/2013
Project start date: 01/11/2010
Duration: 42 months
WP Manager: Emilio Mordini
Revision: 1

Author(s): A.P. Rebera; E. Mordini (CSSC)

Project funded by the European Commission in the 7th Framework Programme (2008-2010)		
Dissemination Level		
PU	Public	Yes
RE	Restricted to a group specified by the consortium (includes Commission Services)	No
CO	Confidential, only for members of the consortium (includes Commission Services)	No

D7.5: Ethical Guidelines

Abstract:

This deliverable provides a set of guidelines for the ethical and legal evaluation of anti-spoofing technologies in biometrics. Perhaps the most significant finding of the Work Package 7 ("Ethical, Social and Cultural Factors") has been the ethical importance of *contextual factors*, i.e. details specific to the implementation of a given biometric/anti-spoofing system (such as the presence or amount of human supervision, the cultural norms of data subjects, and so on). There are always contextual factors so significant that failure to take them into account could potentially lead to miscalculations in the assessment of proportionality of anti-spoofing measures (in terms of potential value gained from anti-spoofing security versus potential threats to fundamental rights). Securing proportionality is an ethical and a legal imperative. It is therefore necessary that anti-spoofing technologies be assessed, as far as possible, with reference to the specific context in which they are to be applied. This is, practically speaking, problematic, for it implies that there cannot, on the whole, be hard and fast principles unquestionably applicable to every case. Accordingly, the guidelines herein presented describe standards for the deployment of anti-spoofing technologies in part by proposing effective ways of identifying, assessing, and reacting to salient contextual variables.

Summary

This deliverable provides guidelines for the ethical and legal evaluation of anti-spoofing technologies in biometrics. The guidelines are based on research in WP7, "Ethical, Social and Cultural Factors".

The guidelines are organised according to **"deployment phases"**. The user is able to consult the guidelines with specific reference to whichever deployment phase is, at the time of consultation, most relevant.

- *The Pre-deployment Phase*: The period during which the need to deploy anti-spoofing technology is identified, and the particular anti-spoofing technology is chosen or developed.
- *The Initial Deployment Phase*: The period in which an anti-spoofing technology is first deployed.
- *The Long-term Deployment Phase*: The period at which the anti-spoofing technology is regularly or routinely used; and (if applicable) the period at which the system is taken out of use or replaced.

We identify **six stakeholder groups**: developers, vendors, end-users, data subjects, society, and regulators. Given the main aim of the guidelines (i.e. to promote ethically acceptable and legally compliant technologies), they are primarily aimed at developers and end-users. However they are also relevant at the other four stakeholder groups.

Legal/policy and ethical considerations are, at a fundamental level, of different kinds. In this document **the policy and fundamental rights framework is taken as defining the minimum standards against which anti-spoofing technologies should be assessed**. We take it as a further question whether compliance with that legal/policy framework is also ethically sufficient. Thus we sometimes propose guidelines or standards that go beyond the minimum required by the policy framework. This reflects the fact that the guidelines aim at promoting ethical deployment of anti-spoofing technology, not merely "policy-compliant" deployment.

We argue that **the most significant factors in ensuring that anti-spoofing technologies are ethically acceptable are *contextual***. That is, the most significant ethical factors derive from specific features of the particular context in which an anti-spoofing technology is deployed. These factors cannot be fully anticipated or resolved in the abstract, or by developers who are never privy to every contextually salient feature. This implies that there cannot be hard and fast principles unquestionably applicable to every case.

Guidelines should be sufficiently broad as to apply in many situations, yet flexible enough that they can be of practical use in (perhaps very) different contexts. To this end, we not only present guidelines for each deployment phase (§§5.1-5.3), but also provide **guidelines for "*Contextual Factors Analysis*"** (§5.6). The contextual factors analysis is intended to support the identification of salient contextual variables which ought to be taken into consideration. With the contextual factors analysis completed, it should be easier to determine how best to proceed with the deployment of a biometric/anti-spoofing technology in a given scenario.

The deployment phase guidelines make further reference to general **guidelines specific to *Privacy and Data Protection* and *Fundamental Rights***. These are presented in §5.4 and §5.5.

Given the close connection between anti-spoofing (in particular) and biometrics (in general), there may be overlap with existing biometrics guidelines (e.g. from ISO). This is not problematic since our guidelines are not, nor were ever intended to be, in tension with existing provision. These guidelines on spoofing are complementary to existing guidelines on biometrics.

As a final comment, it should of course be considered that we present here suggested guidelines and best-practice: ***compliance with these guidelines does not, in and of itself, imply either immunity from legal obligations, or compliance with legal requirements***. The responsibility to ensure compliance with legal requirements is in no way discharged by compliance with the guidelines here presented.

Contents

1. Introduction

This deliverable provides a set of guidelines for the ethical and legal evaluation of anti-spoofing technologies in biometrics. The guidelines are based upon research in Tabula Rasa Work Package 7 (WP7), "Ethical, Social and Cultural Factors". Research in WP7 has been presented in deliverables D7.1, D7.2, D7.3, and D7.4.

- **D7.1. Environmental Scanning Report**
  An inventory of existing regulations, reports, working papers, academic articles (etc.) addressing ethical, legal and policy implications of systems for spoofing prevention in various contexts and applications.
- **D7.2. Interviews Report**
  A report on 31 elite interviews on biometrics, anti-spoofing, and attendant ethical, cultural and social issues.
- **D7.3. Ethical Monitoring Report (Intermediary)**
  An account of the ethical monitoring of the Tabula Rasa project, including the main ethical concerns raised by the project and the internal monitoring procedures employed.
- **D7.4. Workshop Report**
  A report on the "Spoofing and Anti-Spoofing: the Wider Human Context" workshop (Rome, 10-11 May 2012).

The most important findings of these deliverables are summarised (§§3-4) and translated into guidelines §5) below. Earlier work is discussed and expanded upon in the annexes to this document (§§8-9).

The guidelines are organised according to what we call "deployment phases", i.e. the different phases in the development and deployment of anti-spoofing technologies (or biometric technologies equipped with anti-spoofing technologies). There are three deployment phases: the "pre-deployment phase", the "initial deployment phase" and the "long-term deployment phase".

The user may consult the guidelines with reference to the relevant deployment phase. Each guideline specifies the kinds of anti-spoofing technologies to which it is relevant (human supervision; liveness-detection; multi-modality; inherently robust modalities), as well as the stakeholder groups to whom it applies. Use case examples (derived from Tabula Rasa D2.1) are employed as illustrative examples.

In compiling the guidelines we have consulted guidelines and recommendations from the *International Organization for Standardization* (ISO), the *Article 29 Data Protection Working Party* (the European Union's data protection advisory body), and the *Council of Europe*. In

particular we consulted:

- **International Organization for Standardization (ISO)**
  Information technology – Biometrics – Jurisdictional and societal considerations for commercial applications – Part 1: General guidance (ISO/IEC TR 24714-1:2008).
- **Article 29 Data Protection Working Party**
  Opinion 3/2012 on developments in biometric technologies (00720/12/EN WP193).
- **Council of Europe**
  The need for a global consideration of the human rights implications of biometrics (Parliamentary Assembly Resolution 1797 [2011]).

These documents provide guidelines relating to biometrics in general, as opposed to anti-spoofing technologies in particular. Unavoidably there is some degree of crossover between guidelines relating to biometrics and guidelines relating to anti-spoofing. We have tried in this document, as far as possible, always to speak to anti-spoofing. We have not, however, shied away from issues that undoubtedly concern the two. Data protection is a case in point. In this area, anti-spoofing technologies raise questions of legal requirements and ethical best-practice. These questions are structurally similar to the data protection questions raised by biometric systems in general, but we do not, on such grounds, pass them over. Rather, we address them as fully as necessary, even if in so doing we risk running over ground already well-trodden in other reports (such as the abovementioned ISO and Article 29 documents). To do otherwise would result in less than comprehensive guidelines for anti-spoofing.

There may, then, be some overlap with existing biometrics guidelines. This is not problematic since our guidelines are not, nor were ever intended to be, in tension with existing provision. These guidelines on spoofing are complementary to existing guidelines on biometrics.

As a final comment in this introduction, it should of course be considered that we present here suggested guidelines and best-practice: ***compliance with these guidelines does not, in and of itself, imply either immunity from legal obligations, or compliance with legal requirements***. The responsibility to ensure compliance with legal requirements is in no way discharged by compliance with the guidelines here presented.

2. Methodological Considerations

This section aims at clarifying the role of the guidelines by answering four questions:

1.) What are these ethical guidelines for?
2.) To whom are they addressed?
3.) How might they be used?
4.) What are the advantages of following the guidelines?

2.1. What are the Guidelines For?

Anti-spoofing technologies can be, from an ethical or legal point of view (the two are not the same) acceptable or not. In short, the aim of the guidelines presented in this deliverable is to promote ethically acceptable and legally compliant anti-spoofing technologies.

That legal and ethical considerations are, at a fundamental level, of different kinds is fairly obvious. Although many laws are based on ethical norms (for example to commit murder is both illegal *and* inherently unethical), this is not always the case. Indeed when laws are changed, this is often because they are seen to enshrine unethical practices. We do not say in such cases that what was previously, under the old law, both legal and ethical is now, under the new law, illegal and unethical; rather we say that what was previously legal and now illegal was unethical all along (consider the introduction of laws abolishing slavery for example).[1]

The important point here is that we will set out the European policy and fundamental rights framework, but we will not assume that in doing so we thereby satisfy all *ethical* requirements. In some respects it may well be; in other respects it may be necessary to set out guidelines that go beyond what is technically required (e.g. by imposing demands more strict than those imposed by the existing policy and fundamental rights framework).

Accordingly, we review below the policy and fundamental rights framework (§3) and, separately, review the ethical issues that have emerged in WP7 (§4). These reviews are partly based on completed WP7 research; however they go beyond the results earlier deliverables, updating and extending findings and conclusions as required. The guidelines (§5) are derived from these reviews.

[1] Consider also, however, that this observation carries no necessary implication of *progress*, i.e. *from* legislation which is legal and unethical *to* legislation which is legal and ethical (the introduction of racial laws in Nazi and fascist regimes is a case in point).

2.2. To Whom are the Guidelines Addressed?

For their guidelines on the societally acceptable implementation of biometric systems, ISO identifies the following primary stakeholders (ISO, 2008: v):

- Users (i.e. those who use the results of the biometric data)
- Developers of technical standards
- Subjects (i.e. those who provide a sample of their biometric data)
- Writers of system specifications, system architects and IT designers
- Public policy makers

This list is not ideal. First, since the "Users" group cannot use biometric data independently of societal considerations (ethical and cultural norms, legal and policy frameworks, public opinion, etc.), an additional stakeholder group should be included to reflect these issues. Second, it is not clear that "Developers of technical standards" constitutes a stakeholder group entirely distinct from either "Writers of system specifications, system architects and IT designers" or "Public policy makers". Thirdly, the "Subjects" group is restricted to those who provide their biometric data; but there is surely a need to address the concerns and interests of people and groups living in a society in which biometric systems proliferate, regardless of whether they regularly – or indeed ever – provide biometric data. Fourthly, stakeholders also include the vendors of biometric systems. Sometimes system architects may also be vendors, but this need not be so. Clearly those who make their living or have a business stake in the biometrics industry are relevant stakeholders. Finally, "Public policy makers" are not the only relevant public authorities; regulators (such as data protection authorities) should also be considered.

Adopting then, a somewhat broader definition of "stakeholder", we identify six stakeholder groups.

1. ***Developers***. Researchers, engineers, system architects, (etc.), investigating and developing anti-spoofing technologies (in either an academic or commercial setting).
2. ***Vendors***. Commercial enterprises marketing and selling anti-spoofing technologies.
3. ***End-users***. End-users wishing to acquire and deploy anti-spoofing technologies or biometric technologies equipped with anti-spoofing technologies. These include: military agencies, governmental and civil agencies, commercial entities, and individuals.[2]
4. ***Data subjects***. Individuals who will or may be subject to identification, verification, or profiling by biometric technologies equipped with anti-spoofing technologies.[3]

[2] "End-user" is a term used in various ways. Here we use it only to denote individuals and groups deploying biometric systems—not those who are identified or authenticated by them.

[3] "Data subject" is not an ethically neutral term. We use it here only to distinguish people who are identified by biometric systems from those people who, for whatever reason, happen never to be biometrically identified (the

5. ***Society***.[4] Society as a whole, subsections of society (including individuals), formal and informal civil society groups, etc., i.e. any group or individual who has to live in a society in which there is relatively widespread deployment of biometric technologies equipped with anti-spoofing technologies.
6. ***Regulators***. Policy makers and regulators wishing to oversee, shape, and govern the deployment of biometric technologies equipped with anti-spoofing technologies.

Given the main aim of the guidelines (i.e. to promote ethically acceptable and legally compliant technologies), they are primarily aimed at **developers** and **end-users**. However they are also aimed at the other four stakeholder groups.

We argue below that the most significant factors in ensuring that anti-spoofing technologies are ethically acceptable are *contextual*. That is to say, the most significant ethical factors derive from specific features of the particular context in which an anti-spoofing system is deployed. These factors cannot be fully anticipated or resolved in the abstract, or by developers who are never privy to every contextually salient feature. Nonetheless, the guidelines are intended to support **developers** in developing anti-spoofing technologies which, when deployed, do not *from their general design or functionality* cause or aggravate ethical issues.[5]

Given the impact of particular contexts, the guidelines support **end-users** in:

1. identifying ethically salient features of a use-context;
2. assessing whether the use of anti-spoofing technology is ethically acceptable in a given use-context;
3. ensuring that anti-spoofing technology is deployed in an ethically acceptable manner.
4. continuously monitoring the ethical acceptability of a given deployment of an anti-spoofing technology.

The guidelines speak directly to developers and end-users, but should nevertheless be of interest to all the other stakeholder groups. Since vendors' customers—i.e. end-users—need anti-spoofing technologies that can be easily deployed in an ethically acceptable manner, **vendors** should demand that the systems they source from developers do not, from their design or functionality,

latter group fall within stakeholder group 5, "Society").

[4] We acknowledge that there is some difficulty in including so large (and, frankly, *vague*) a stakeholder group. Our intention is to simply to signal that the use of anti-spoofing technology has implications for a society's way of life. To mitigate the difficulty we have indicated clearly above that by "society" we recognise not only society *as a whole* but also specific groups within it (civil society groups, individuals, etc.).

[5] One ethical issue concerns the conceptualisation of spoofing as deceptive and (therefore) illegitimate. Might there be contexts in which spoofing is acceptable? Might there be contexts in which spoofing is actually an ethical imperative (resistance to an illegitimate authority, perhaps)? In fact, spoofing raises a theoretical problem similar to that of lying (cf. Rebera et al, 2013). However this problem is of less immediate importance in the present document, since in Tabula Rasa we assume legitimacy of contexts (i.e. that spoofing is (rightly) illegal and detecting it is always desirable).

cause or aggravate ethical issues. Vendors may therefore find the guidelines interesting as a source of information on what they should expect from systems. They may also find the guidelines of use in determining how much information they should share with their customers (i.e. end-users).

Data subjects, i.e. the people who are going to be identified or authenticated by biometric systems equipped with anti-spoofing technologies, should be aware of their rights and of what standards they may expect from end-users. The guidelines may be of interest to them insofar as they describe those standards. The same may be said of **society**. Specific civil society organisations advocating privacy, civil liberties, etc., may take an interest in the guidelines, and may wish to either endorse or suggest improvements to them.

Finally **regulators**, specifically policy makers, may take an interest in the guidelines insofar as they propose measures which go beyond the current policy and fundamental rights framework.[6] Regulators may also find the guidelines of use insofar as they promote transparency with respect to, say, proportionality (since the policy framework legislates for proportionality, e.g. in the processing of personal data).

2.3. How Might the Guidelines be Used?

The guidelines are presented in such a way that they can be fruitfully consulted at the different "deployment phases" in the lifetime of a biometric system equipped with anti-spoofing technology.

ISO lists the following stages in the life cycle of a biometric system (ISO, 2008: 1):

1. Capture and design of initial requirements, including legal frameworks;
2. Development and deployment;
3. Operations, including enrolment and subsequent usage;
4. Interrelationships with other systems;
5. Related data storage and security of data;
6. Data updates and maintenance;
7. Training and awareness;
8. System evaluation and audit;
9. Controlled system expiration.

For simplicity of presentation we condense these stages into three broad deployment stages around which our guidelines are organised: the "pre-deployment phase", the "initial deployment

[6] Of course it should be noted that it is not always necessary or desirable to regulate to ensure ethical practice. Hence where the guidelines suggest measures more strict than the framework this ought not to be immediately taken as a indicating a "policy gap". (For example, the lack of regulation to ensure interpersonal politeness is no policy gap: regulating politeness would be counterproductive.)

phase", and the "long-term deployment phase". The user is able to consult the guidelines with specific reference to whichever deployment phase is, at the time of consultation, most relevant.

- **The Pre-deployment Phase**
  The period during which the need to deploy anti-spoofing technology is identified, and the particular system/ anti-spoofing technology is chosen or developed.
- **The Initial Deployment Phase**
  The period in which an anti-spoofing technology is first deployed.
- **The Long-term Deployment Phase**
  The period at which the anti-spoofing technology is regularly or routinely used; and (if applicable) the period at which the system is taken out of use or replaced.

Of course there are several ways in which guidelines of this kind could reasonably be presented. The "deployment phase" approach seems to us the most appropriate. We could, for example, have structured the guidelines around *individual biometric modalities*. But this proposal is suboptimal because both the vulnerability of a system focused on a given modality *and* the ethical issues arising in relation to that modality are subject to some degree of variation according to the context in which the system is deployed.

Noting this, we could instead have provided guidelines for a *representative selection of use cases*. This approach is perhaps more appropriate than the modalities approach, since it superficially sidesteps a greater proportion of contextual variability. On reflection however, the benefit of this approach—that it speaks to relatively concrete example scenarios—is offset by the difficulty that remains to be faced in generalising from a limited set of concrete examples. A better approach is to avoid relying on particular example use-cases, and instead to identify the kinds of contextual variables that are likely to be common across all use-case scenarios. The "deployment phase" approach allows us to do this.

Finally, we could plausibly have structured the guidelines around the different *categories of anti-spoofing technology*. Thus we might have presented "guidelines for the use of liveness-detection", "guidelines for the use of inherently robust modalities", and so on. Of the organisational principles we have declined, this is arguably the best. Despite this we have rejected it as the *primary* organisational principle on the following grounds. To set out guidelines for the ethical deployment of (say) liveness-detection technologies seems to presuppose a positive answer to the question: "is liveness-detection technology an appropriate measure in the class of use-case scenarios to which the liveness-detection guidelines refer?" Yet this is a serious ethical (and in some cases legal) question, to which it would be problematic to simply presuppose a positive answer. Of course adopting the different categories of anti-spoofing technology as the primary organisational principle would not *rule out* addressing that question at some point. But it is rather more elegant and clear to set out the guidelines in such a way that there is a dedicated section (i.e.

the "pre-deployment phase") in which the question of appropriateness and proportionality naturally arises as a central issue.

The attraction of the "deployment phase" approach is clear. Fortunately, it in no way rules out our speaking to each of the categories of anti-spoofing technology *within* each deployment phase. So that is the course we have steered: for each set of guidelines, it has been made clear to which categories of anti-spoofing technology it relates.

Finally, to ensure that the guidelines and the principles that support them are sufficiently clear – and in continuance of a running Tabula Rasa theme – the guidelines section makes reference to illustrative example use-cases, based on those set out in D2.1 (and further discussed in D7.1, Appendix §2).[7]

2.4. What are the Advantages of Following the Guidelines?

Benefits to be gained from adhering to the guidelines include the following, suggested by ISO (ISO, 2008: v), with respect to biometrics in general:

- Enhanced acceptance of biometrics systems by subjects;
- Improved public perception and understanding of well-designed systems;
- Smoother introduction and operation of systems;
- Potential long-term cost reduction;
- Increased awareness of accessibility-related issues;
- Adoption of commonly approved good privacy practice.

Adapting and extending, the table below (Table 1) presents the following possible advantages for each stakeholder group.

[7] Note that the use-case examples are used as only to illustrate and contextualise salient ethical issues—the guidelines are not *derived from* the use-cases.

Potential Advantage		Beneficiary Stakeholder	Principal Benefit to Stakeholder
1	Early identification of ethical problems arising from system design/functionality.	Developers	May enable a more efficient allocation of resources (funding, time, etc.), with the subsequent benefit of being able to devote more resources to technologies likely to be demanded by vendors and end-users.
		Vendors	May make it easier to source legally, ethically, societally acceptable (and hence *marketable*) technologies.
		End-users	May make it easier to source legally, ethically, societally acceptable (and hence *easily deployable*) technologies.
2	Promote common understandings of ethical issues arising in relation to the deployment of anti-spoofing technologies.	Developers	May promote the development of *ethics-*, *data-protection-*, *privacy-* (etc.) *by-design* solutions.
		Vendors	May make it easier to identify and market optimal solutions to potential problems.
		End-users	In anticipating potential issues, end-users may find it easier to identify the optimal anti-spoofing technologies for their specific needs.
		Data subjects	Greater awareness of ethical issues on the part of end-users. May find potential issues resolved in advance.
		Society	May enjoy a more acceptable/less controversial deployment of biometric systems/anti-spoofing technologies.
		Regulators	May see a greater compliance with requirements and standards.
3	Compliance with legal/policy frameworks, ethical and cultural norms, etc. may promote societal understanding and acceptance of biometrics/anti-spoofing.	Developers	May serve to safeguard research funding and/or the market for biometric anti-spoofing technologies.
		Vendors	May maintain/grow the market for anti-spoofing technologies.
		End-users	Greater societal understanding and acceptance may reduce data subjects' dissatisfaction, increasing the efficiency the end-users' deployment of the technology. Greater efficiency, reduced complaints, etc. may have financial or other benefits.
		Data subjects	May find biometric/anti-spoofing technologies more acceptable and user-friendly.
		Society	May find biometric/anti-spoofing technologies more acceptable and user-friendly.
		Regulators	May receive fewer complaints if the technology is more societally acceptable.
4	Encourages commonly accepted standards for ethical best-practice with respect to anti-spoofing technologies.	Developers	May simplify development procedures (through standardisation).
		End-users	Encourage openness with respect to privacy, ethics (etc.) policies, building trust with data subjects and society
		Data subjects	May make it easier to identify and distinguish acceptable and unacceptable deployment by end-users. Promotes active protection of fundamental rights (e.g. dignity, privacy, data protection).
		Society	May promote greater awareness of the importance of fundamental rights and ethics in the development of technology and the information society in general.
		Regulators	May make it easier to see where legal/policy frameworks are most/least effective, and to identify gaps in policy or technology-specific/industry standards.

Table 1. Possible advantages for stakeholder groups

2.5. Summary

- The guidelines are organised according to three “deployment phases”. They can be fruitfully consulted with reference to whichever deployment phase is, at the time of consultation, most appropriate.
- The “deployment phase” approach is appropriate because the most significant ethical factors derive from specific features of the particular context in which a biometric system is deployed.
- For each deployment phase, the guidelines are linked to specific stakeholders and categories of anti-spoofing technology.
- The guidelines include discussion of illustrative example use-cases, based on those originally set out and discussed in D2.1 and D7.1.
- Following the guidelines may result in a number of advantages for the key stakeholders.

3. European Policy and Fundamental Rights Framework

This section summarises the main aspects of the European policy and fundamental rights framework relevant to the guidelines. Rather than offer a lengthy discussion updating issues addressed in earlier WP7 research – which would distract from the primary purpose of the present document – we present here a summary. Detailed discussion is to be found in Annex 1 (§8) below.

The European legal framework concerning biometrics – and so spoofing – is firmly based around data protection and fundamental rights legislation, in particular the *European Data Protection Directive* (95/46/EC), the *European Convention on Human Rights* (ECHR), and the *Charter of Fundamental Rights of the European Union* (CFREU).

3.1. Biometric Data = Personal Data? Biometric Data = Sensitive Data?

The Data Protection Directive (95/46/EC) governs the processing of personal data.[8] "Sensitive data" is a special subcategory of personal data.[9] In D7.3 we concluded that biometric data is *always* personal data, and *sometime* sensitive.

For present purposes, the following points should be noted.

- Several data protection concepts or terms of art are unclearly defined. These include "personal data", "sensitive data", "identifiable natural person", and "health-related data".
- Biometric data (including templates) should be considered as *personal data*, and handled – at a minimum – in accordance with data protection regulations on personal data.
- In many cases, biometric data should be considered *sensitive*, and processed as such.

[8] "Personal data" and "processing" are defined thus: "*For the purposes of this Directive […] 'personal data' shall mean any information relating to an identified or identifiable natural person ("data subject"); an identifiable person is one who can be identified, directly or indirectly, in particular by reference to an identification number or to one or more factors specific to his physical, physiological, mental, economic, cultural or social identity*" (95/46/EC, Art. 2). "*'[P]rocessing of personal data' ('processing') shall mean any operation or set of operations which is performed upon personal data, whether or not by automatic means, such as collection, recording, organization, storage, adaptation or alteration, retrieval, consultation, use, disclosure by transmission, dissemination or otherwise making available, alignment or combination, blocking, erasure or destruction*" (95/46/EC, Art. 2).

[9] Sensitive data is "[P]ersonal data revealing [of] racial or ethnic origin, political opinions, religious or philosophical beliefs, trade-union membership, and […] data concerning health or sex life" (95/46/EC, Art. 8).

- The proposed Data Protection Regulation includes a definition of "data concerning health" (EC, 2012: 42). This fails, however, to resolve the question of whether biometric data "concerns health".
- From an ethical (as opposed to purely data protection) perspective, biometric data should be handled with caution, especially – but not only – if it is (from a data protection perspective) sensitive.

3.2. Proportionality

Proportionality is one of the most important data protection principles. Its requirement is captured in 95/46/EC by Article 6 (1b), which states that data must be "adequate, relevant and not excessive in relation to the purposes for which they are collected and/or further processed".

Essentially, the principle of proportionality enjoins the securing of a satisfactory balance between (i) the level of data collected and processed and (ii) the importance or value of the purpose for which they are collected and processed. However proportionality – what it means and how it can be ensured – is not perfectly clear.

The main points to be noted are these:

- Proportionality is one of the most important data protection principles.
- The processing of personal data should be a necessary means to a legitimate end.
- A number of factors must be considered in the balance of proportionality, several of which are contextually variant.
- Since contexts are not static, proportionality should be regularly reassessed.
- The proportionality of anti-spoofing provision should be judged in conjunction with the proportionality of the system as a whole: enhancing system security is merely a sub-element of the wider goal of ensuring an efficient and effective biometric system.

3.3. Consent

95/46/EC defines *consent* as "any freely given specific and informed indication of his wishes by which the data subject signifies his agreement to personal data relating to him being processed" (Art. 2(h)). However a European Commission Communication of 2010 reports a number of concerns regarding the concept of *consent* (EC, 2010).

Since anti-spoofing provision may involve the processing of personal data, and since it may sometimes involve the processing of *sensitive* personal data – a form of processing that may require "explicit consent" (95/46/EC, Art 8(2a)) – it is necessary to examine the concept of consent more closely. The following points draw out the implications of the approach applied in

95/46/EC. (The specific relevance of these implications to the deployment of anti-spoofing technologies is discussed at length in Annex 1.)

- There must genuinely be more than one option: to consent or not to consent.
- Both options, to consent or not to consent, must be realistically possible for the subject, and neither should be disproportionately (dis)advantageous.
- The subject must be genuinely aware that there is a decision to be taken (to consent or not to consent) and that the decision is their own.
- A subject's stated decision to (not) consent must be their own. The decision need not be entirely independent of outside influence, but it should not be subject to undue influence or coercion**.**
- In order to be in a position to give consent, a data subject should understand (or have been given sufficient opportunity to come to understand) the nature and consequences of the action or data processing to which they are potentially subject.
- The data processor's means of recognising consent should be clear to the data subject. The danger of the subject "inadvertently" giving consent should be minimised.[10]
- There should be a record of the giving of explicit consent (e.g. in writing).

[10] The proposed Data Protection Regulation (EC, 2012) makes clear that consent should be *affirmative* (i.e. *actively* signaled by the subject).

4. Ethical Issues

In this section we summarise the ethical issues arising from the development and deployment of anti-spoofing technologies. More detailed discussion building on earlier WP7 work is presented in Annex 2 (§9)

The subsection is structured around three thematic areas:

- Data protection, privacy, and fundamental rights;
- Spoofing, deception, and honesty;
- Trust, transparency, and secrecy;

4.1. Data protection, privacy, and fundamental rights

For the most part, these issues are summarised above (§3), and discussed at length in Annex 1 (§8). Here, however, we approach the topic from a more strictly ethical (less policy-specific) direction.

- One reason behind the need for strong data protection is that personal data is increasingly economically valuable. That there is a market for personal information is broadly understood, yet the role of individual data subjects in that marketplace remains unclear. There is no consistently applied means of assessing the value of personal data.
- Anonymisation and data minimisation should be pursued whenever possible. The less data a system gathers and retains, the wider its likely range of acceptable deployment scenarios.
- Data minimisation is also important because it reduces the risk of function creep.
- Understanding of the current vulnerabilities of biometric systems is imperfect. We have also extremely imperfect knowledge of the *future* vulnerabilities of biometric systems.
- The gathering of data by biometric/anti-spoofing systems is of concern with respect to fundamental rights beyond data protection. These include dignity and integrity.[11]
- Dignity can be threatened by the gathering of *too much* personal information about an individual.
- Dignity can be threatened if the data gathered is *sensitive* (in the technical data protection sense) or *intimate* (i.e. data which is, in some sense, extremely direct, and which gets to the core of who and how one fundamentally is; e.g. information about one's personality or current emotional state).
- One's sense of identity can be imposed, influenced, or controlled by others. Anti-spoofing technology should not contribute to the imposition of negative self-

[11] On dignity and integrity see Tabula Rasa D7.1 (§§4.4, 4.3 respectively).

conceptions on data subjects (e.g. through a supposed connection between anti-spoofing and deviancy or suspicion).

- People have a legitimate interest in anonymity in various circumstances: covert or low-visibility gathering of data can compromise this interest. (No article of the European fundamental rights framework specifically sets out a right to anonymity; however it arguably falls somewhat within the scope of other provisions, including the rights to privacy, dignity, and integrity.)

4.2. Spoofing, deception, and honesty

The supposed conceptual link between spoofing and deviancy is stronger than that supposed between biometrics and deviancy.

- While spoofing certainly is deceptive, it does not follow that it is always morally impermissible (Rebera et al, 2013).
- It is necessary to consider the question of *who spoofs and why?* Or to put it another way: *under what circumstances could it be morally justifiable to spoof a biometric system*? Since, in general, spoofing is justifiable only when something has gone wrong – i.e. when the spoofer's legitimate interests have been compromised – answering this question is likely to be a key step in the development of the guidelines.
- In the European context, the justifiability question boils down to one of proportionality: under what circumstances could a legitimate authority's deployment of a biometric system equipped with anti-spoofing provision be ethically unacceptable? The broad answer is: when it infringes, for no good reason, people's fundamental rights and legitimate interests.
- Actors deploying biometric systems should be honest about their reasons for so doing and the details of the deployment. Their honesty should be *active*, i.e. they should actively take steps to inform relevant stakeholders. Active honesty does not entail broadcasting details to *everyone*. Rather, it entails informing *appropriate* stakeholders.

4.3. Trust, transparency, and secrecy

Leading on from openness and active honesty, we turn now to transparency and secrecy.

- In advocating against secrecy, what is promoted is a culture or ethos whereby intentional concealment is the exception rather than the rule (we do not deny that secrecy may be necessary in certain circumstances). The intention should be to bring into the light as much information as is possible, bearing in mind security considerations.
- In promoting transparency, what is promoted is the clearing away of obstructions barring from view the structures by and through which anti-spoofing technologies are developed and deployed. The intention here is to make what is in the light clearer to view. However, *transparency* does not entail *complete publicity*.

- Public trust and confidence in biometrics, anti-spoofing, and ICT in general, needs to be encouraged by developing and publicly promoting secure systems.
- Openness should apply to technologies *and* operational procedures.
- Information is only helpful if its target audience can understand it. Setting out technical details in comprehensible form is a significant challenge.
- Vendors of biometric systems may have an economic interest in not revealing spoofing vulnerabilities. Overcoming this reluctance may be possible through (e.g.) legislation, industry standards and codes of conduct, or economic incentives.
- Vendors should make available to customers sufficient information regarding security and vulnerabilities that the customer is able to make a well-informed judgment as to whether the system in question is a good fit for the specific context(s) in which they intend to deploy it. This suggests a shared responsibility between vendor and customer: the customer must make an accurate analysis of the deployment context they are considering; the vendor must provide sufficient information and in a comprehensible form.
- Other things being equal, vendors should be completely open about the spoofing vulnerabilities of the systems they sell. (There is no need to actively publicise vulnerabilities; but they should not be hidden either.)
- If releasing information concerning the vulnerabilities of pre-deployment stage technologies would compromise those technologies' likelihoods of reaching the deployment stage, that information may be held confidential.

5. Ethical Guidelines

This section sets out guidelines for the ethical and legal evaluation of anti-spoofing technologies in biometrics. WP7's research has shown the ethical importance of **contextual factors** (such as the presence or amount of human supervision, the cultural norms of data subjects using the system, and so on). There are always contextual factors so significant that failure to take them into account could lead to mistaken assessments of vulnerabilities, and hence of what constitutes a proportional response to those vulnerabilities. In order to secure proportionality—which is both an ethical and a legal imperative—anti-spoofing technologies have to be assessed with reference to their specific deployment contexts. This implies that there cannot be hard and fast principles unquestionably applicable to every case.

The guidelines are organised by deployment phases. The rationale for this is explained above (§2). The tables presenting the guidelines also indicate the stakeholders to whom, and anti-spoofing techniques to which, each guideline applies.

The deployment phase guidelines make reference to general guidelines specific to *Privacy and Data Protection*, *Fundamental Rights*, and to a "*Contextual Factors Analysis*". These are presented in §5.4, §5.5 and §5.6 respectively. The Contextual Factors Analysis is intended to support the identification of salient contextual factors which ought to be taken into consideration.

A useful starting point for thinking about contexts is Nissenbaum's (2010) well-received notion of "contextual privacy". As far as informational privacy is concerned, Nissenbaum's approach provides useful guidance for understanding what kinds of parameters must be considered when assessing the salient variables of a context. We need not adopt Nissenbaum's particular account of privacy, but her discussion of contextually salient factors is useful. However, two points should be noted. First, Nissenbaum's parameters must be adapted to present needs. Second, since our remit extends wider than informational privacy, further variables remain to be identified.[12]

[12] For Nissenbaum, common intuitions concerning what are and are not acceptable informational exchanges are governed by "informational norms". These vary according to the precise calibration of four parameters (Nissenbaum, 2010, 140-147): "Contexts" (i.e. broad background conditions determining general roles that actors play, kinds of activities carried out, etc.); "Actors" (i.e. the parties to an informational exchange); "Information Types" (i.e. what the information is about); and "Transmission Principles" (i.e. overarching principles constraining the sharing of information). (To give a brief example, imagine two doctors discussing a patient. In a medical context (**context**), two doctors may discuss a patient's (**actors**) medical condition (**information type**), respecting confidentiality norms (**transmission principles**). While this situation may respect the patient's privacy, changing any of the parameters may results in a privacy violation. For instance, two doctor's discussing a patient *while at their children's school sports-day* (**change of context**), is unacceptable; a doctor discussing *with her husband* a patient (**change of actor**) is unacceptable; a doctor revealing to another doctor (for no medical reason) *a patient's criminal record* (**change of information type**) is unacceptable; two doctors discussing a patient *in front of other patients* (**violation of transmission principle**) is unacceptable.).

Finally, it should of course be considered that we present here suggested guidelines and best-practice: ***compliance with these guidelines does not, in and of itself, imply either immunity from legal obligations, or compliance with legal requirements***. The responsibility to ensure compliance with legal requirements is in no way discharged by compliance with the guidelines here presented.

5.1. The Pre-deployment Phase

The pre-deployment phase is the period during which the need to deploy anti-spoofing technology is identified, and the particular anti-spoofing technology is chosen or developed.

#	Guideline	Directly Relevant to which Stakeholders?	Relevant to which Countermeasures?
PD.1	Developers should – as far as possible – consider the kinds of context in which the technologies they develop are likely to be deployed. Broad use contexts could include the following: **(a) Physical Access Control**: Use of biometrics to identify data subjects accessing a physical space (room, building, etc.). **(b) Logical Access Control**: Use of biometrics to identify data subjects accessing a computer, network, etc. **(c) Civil/Civic Identification**: Use of biometrics, by government authorities, to identify data subjects for official purposes. **(d) Criminal Investigation/Security Identification**: Use of biometrics to identify suspects, people arrested, monitoring and surveillance uses, etc. **(e) Commercial Authentication**: Use of biometrics to identify data subjects as consumers, perhaps remotely, for commercial purposes	- Developers	- Liveness Detection - Multi-modality - Inherently Robust Modalities

	(including marketing). As far as possible, careful consideration should be given to the kinds of ethical issues that these contexts might raise. For example, would an anti-spoofing system which requires active (and conscious) data subject participation (e.g. challenge-response methods) be more or less appropriate in a given context? *Specific* details of contexts will likely be decisive – and of these developers are probably unaware – but by considering *broad* contexts, developers can contribute to the development of anti-spoofing technologies which respect privacy and fundamental rights.		
PD.2	Developers should identify relevant stakeholders (end-users, likely data subjects, etc.). If any stakeholder group is likely to be *significantly affected* – positively or negatively – with respect to ethics and fundamental rights, they should be consulted in order to identify possible privacy, data protection, or fundamental rights issues. If possible, technological fixes or mitigation strategies can be designed-in to the technology. Note that detailed stakeholder analysis may not be necessary in all cases, but should always be an option.	- Developers	- Liveness Detection - Multi-modality - Inherently Robust Modalities
PD.3	Measures to protect personal data and privacy should be built-in to technologies. Privacy-enhancing features should be turned on by default. Relevant measures may include the following. (See also "Guidelines on Privacy and Data Protection" (§5.4 below).) **(a) Avoid Identifiers**: Data collected for purposes of anti-spoofing should, if at all possible, not identify the data subject. Other things being equal, the least intrusive and the least identifying options should always be favoured.	- Developers	- Liveness Detection - Multi-modality - Inherently Robust Modalities

	(b) Personal Data Minimisation: The collection of redundant data[13] should be prevented; if it is collected, measures to ensure its quick and easy erasure should be available. **(c) Personal Data Anonymisation**: If possible, data anonymisation by default should be built-in as an option. In any case, wherever possible, data should be anonymised. **(d) Personal Data Security**: All reasonable technological measures to ensure the integrity of data processed by a technology should be built-in. **(e) Eliminating Health Data**: All reasonable steps should be taken to ensure that no data relating to health is captured. ***Liveness detection*** is likely to use data relating to health. If it is genuinely necessary to use such data, it should be used and then immediately erased.		
PD.4	Where possible, anti-spoofing features should be built-in to systems in a modular fashion. That is to say: it should be possible to turn on/off single elements on demand and separately, in order to allow data subjects to accept/refuse only *part* of the system. Systems lacking this capacity will be, in effect, "all or nothing" systems, with little or no possibility for contextually-driven flexibility.	- Developers	- Human Supervision - Liveness Detection - Multi-modality - Inherently Robust Modalities
PD.5	On no account should any aspect of the user's interaction with the technology cause them any harm or affront to their fundamental rights. They should never be humiliated, embarrassed, or made to feel inadequate. The system's default configuration should be the most fundamental-rights-	- Developers - Vendors - End-users	- Liveness Detection - Multi-modality - Inherently Robust Modalities

[13] Data is redundant if the system or process for which it is gathered would function or run tolerably well in its absence.

	enhancing. The following factors should be considered in advance of the deployment of anti-spoofing technology. See also "Guidelines on Fundamental Rights" (§5.5 below). **(a) Dignity**: No aspect of the user's interaction with the technology should cause any affront to their dignity. E.g. there should be no requirement to reveal intimate parts of the body. **(b) Integrity**: No aspect of the user's interaction with the technology should cause any affront to their physical or psychological integrity. The user should suffer no harm whatsoever in using the technology. **(c) Non-Discrimination**: The user's interaction with the technology should not involve or lead to illegitimate discrimination. For example, no particular group or community should be disproportionately disadvantaged by the technology. If, e.g., a group is unable (or significantly less able) to use the technology, alternatives should be available to ensure their access to whatever the technology manages access to (however these alternatives may be provided by end-users in deployment).		
PD.6	Prior to deploying a system, a data protection impact assessment should be carried out (or commissioned) by the end-user.[14] The data protection impact assessment should be conducted alongside a "Contextual Factors Analysis".[15] The data protection impact assessment should comply with any local practices, such as standards or requirements from national or local supervisory bodies (for an overview of major issues, see "Guidelines on Privacy and Data Protection", §5.4).	- End-users	- Human Supervision - Liveness Detection - Multi-modality - Inherently Robust Modalities

[14] On "impact assessments" see Annex 1 (§8.4).
[15] See the "Guidelines on Contextual Factors", §5.6.

PD.7	Prior to deploying a system, a privacy impact assessment should be carried out (or commissioned) by the end-user.[16] The privacy impact assessment should be conducted alongside a "Contextual Factors Analysis".[17] The privacy impact assessment should comply with any local practices, such as standards or requirements from national or local supervisory bodies (for an overview of major issues, see "Guidelines on Privacy and Data Protection", §5.4).	- End-users	- Human Supervision - Liveness Detection - Multi-modality - Inherently Robust Modalities
PD.8	Prior to deploying a system, an ethical impact assessment should be carried out (or commissioned) by the end-user.[18] The ethical impact assessment should be conducted alongside a "Contextual Factors Analysis".[19] The ethical impact assessment should comply with any local practices, such as standards or requirements from national or local supervisory bodies (for an overview of major issues, see "Guidelines on Fundamental Rights", §5.5).	- End-users	- Human Supervision - Liveness Detection - Multi-modality - Inherently Robust Modalities
PD.9	Prior to deploying a system, an examination of the relevant contextual factors – a "Contextual Factors Analysis"[20] – should be conducted (this may be part of the data protection, privacy, or ethical impact assessment; or it may be a separate exercise). (See "Contextual Factors Analysis", §5.6.)	- End-users	- Human Supervision - Liveness Detection - Multi-modality - Inherently Robust Modalities
PD.10	Once all analyses and impact assessments are complete, end-users should produce a clear and unambiguous written statement of their plans. This should include at least: **(a)** The purpose of the biometric/anti-spoofing system.	- End-users	- Human Supervision - Liveness Detection - Multi-modality - Inherently Robust Modalities

[16] On "impact assessments" see Annex 1 (§8.4). *Privacy* impact assessment and *data protection* impact assessment are sometimes held apart as separate enterprises (e.g. De Hert 2012). We keep them distinct in these guidelines but accept that they could be conducted together, as a single integrated activity.
[17] See the "Guidelines on Contextual Factors", §5.6.
[18] On "impact assessments" see Annex 1 (§8.4).
[19] See the "Guidelines on Contextual Factors", §5.6.
[20] See the "Guidelines on Contextual Factors", §5.6.

	(b) The kind of biometrics/anti-spoofing techniques used. **(c)** Justification of the kind of biometrics/anti-spoofing techniques used (including explanation of why less intrusive/invasive measures would not be sufficient). **(d)** Whether (and how) the anti-spoofing system will interact with any other biometric or identification systems in the proposed deployment context. **(e)** Statement of relevant stakeholders and the results of any consultation to have taken place **(f)** List of relevant ethical factors (including, e.g., function creep) – including those specific to the given context. **(g)** List of mitigation strategies for each relevant ethical factor. **(h)** Notice of any official approvals sought (e.g. from data protection authorities). **(i)** List of "Critical Incidents". These are incidents so grave that, should they occur, the whole system should be immediately stopped, and re-assessed (cf. ID.7).		
PD.11	If the technology proposed entails any of the following, then the local data protection authority should be informed. Any advice they offer should be closely followed.	- End-users	- Human Supervision - Liveness Detection - Multi-modality - Inherently Robust

	(a) the processing of large amounts of personal data; **(b)** the processing of any amount of sensitive data; **(c)** the processing of personal data in the absence of the data-subjects' explicit consent; **(d)** the gathering of personal data in public spaces;		Modalities
PD.12	Although developers should not lie about the capabilities of emerging anti-spoofing techniques, they are not under an obligation to publicise vulnerabilities or problems with technologies that are not yet in use or on the market.	- Developers	- Liveness Detection - Multi-modality - Inherently Robust Modalities
PD.13	Vendors should be honest, open, and accurate regarding the vulnerabilities of the systems they sell. They may require support in this regard from consultation with developers (on technical capacities), end-users (on operational or procedural factors), and regulators (on industry or legal standards).	- Vendors - Developers - End-users - Regulators	- Human Supervision - Liveness Detection - Multi-modality - Inherently Robust Modalities
PD.14	Efforts should be made to promote standardisation, a common vocabulary, and as robust a measure of vulnerabilities as is possible.[21] If necessary, the use of harmonised standards and vocabularies should be incentivised by competent authorities (e.g. policy makers).	- Regulators - Developers - Vendors - End-users	- Human Supervision - Liveness Detection - Multi-modality - Inherently Robust Modalities

Table 2. Guidelines for the pre-deployment phase.

[21] Due to the unpredictable influence of contextual factors, precise measures of vulnerabilities are not possible. But this does not mean that a harmonised approach – i.e. one which enables a meaningful comparison of different systems within a single context – is not possible.

5.2. The Initial Deployment Phase

The initial deployment phase is the period in which an anti-spoofing technology is first deployed and begins to be regularly or routinely used.

#	Guideline	Directly Relevant to which Stakeholders?	Relevant to which Countermeasures?
ID.1	The anti-spoofing technology should be deployed in accordance with the findings of any impact assessments (data protection, privacy, ethical), as well as in accordance with any recommendations from data protection authorities.	- End-users	- Human Supervision - Liveness Detection - Multi-modality - Inherently Robust Modalities
ID.2	The anti-spoofing technology should be deployed in accordance with the findings of the "Contextual Factors Analysis". The Contextual Factors Analysis will highlight the main variables and issues that must be taken in to account when identifying the relevant ethical factors.	- End-users	- Human Supervision - Liveness Detection - Multi-modality - Inherently Robust Modalities
ID.3	Any privacy-enhancing settings or configurations that the anti-spoofing technology offers should be activated. It should not be assumed that the optimal configuration is set up as standard. Where possible, the technology (and any discrete subparts) should be "opt-in" for users (i.e. they should actively signal consent).	- End-users	- Liveness Detection - Multi-modality - Inherently Robust Modalities
ID.4	Staff overseeing the system (if any) should be well-trained in its use, but must also be able handle data subjects sensitively (e.g. to explain to them why the system is necessary, whether there are any implications, what are their modes of recourse in the case of complaint, and so on). In general, it should be considered that operating procedures are not independent of more "technical" aspects of a system, and should, therefore, be informed by all impact assessment findings, and by any adopted standards (e.g. "privacy by design").	- End-users	- Human Supervision - Liveness Detection - Multi-modality - Inherently Robust Modalities

	Staff may require special training in dealing with users having difficulty using the technology (e.g. children, people with disabilities).		
ID.5	Measures should be taken to ensure that data subjects may have, if they need it, access to clear and concise information on the system processing their data, and on their rights. Such information may be provided verbally, in writing, on a website or by email. It is best provided by the end-users, as they are familiar with relevant contextual factors. But it may be largely based on information or guidelines provided by regulators or industry groups, if such resources exist. A named person responsible should be appointed, in order that data subjects can seek redress if necessary.	- End-users - Regulators	- Human Supervision - Liveness Detection - Multi-modality - Inherently Robust Modalities
ID.6	Relevant industry standards – whether from the biometrics industry, or from the industry in which a biometric system is deployed – for health and safety, privacy, data protection, and the protection of fundamental rights should be consulted and followed.	- End-users	- Human Supervision - Liveness Detection - Multi-modality - Inherently Robust Modalities
ID.7	A list of "critical incidents" should be developed (cf. PD.10) These are incidents so grave that, should they occur, it would follow that something in the system design or implementation was seriously amiss. These incidents should thus never occur, but if they do the whole system should be ***immediately stopped***, and re-assessed. Critical incidents might include: - Physical or psychological harm to a data subject; - Failure to adhere to a legal requirement (e.g. an aspect of data protection law); - Illegitimate discrimination against any individual. This list is obviously incomplete. It should be compiled with reference to contextual factors.	- End-users	- Human Supervision - Liveness Detection - Multi-modality - Inherently Robust Modalities
ID.8	A short time after the anti-spoofing technology is deployed, a review should be conducted to identify any possibilities to better protect privacy and fundamental rights. It should be ensured that the system and processes are functioning reliably and as expected. The review should ensure that all reasonable steps have been taken, and that all guidelines, industry standards, and data protection authority recommendations have been adhered to. If further possible steps are identified, they should be implemented, and then the review	- End-users - Regulators - Data subjects - Society	- Human Supervision - Liveness Detection - Multi-modality - Inherently Robust Modalities

	process revisited shortly after. Steps should be taken to assess the acceptability of the system to people using it. Any problems identified should be swiftly and effectively dealt with. Should any serious problems be identified, the end-user should consult regulators (e.g. data protection authorities) or civil society groups (e.g. privacy advocacy groups) to establish the most suitable response. Should any “critical incidents” occur the system should be suspended immediately (cf. ID.7)		

Table 3. Guidelines for the initial deployment phase.

5.3.The Long-term Deployment Phase

At the long-term deployment phase the biometric system with anti-spoofing technology is established and regularly or routinely, used: staff (overseers, administrators, supervisors, etc.) are familiar with the technology; biometric subjects are accustomed to being identified by the system.

#	Guideline	Directly Relevant to which Stakeholders?	Relevant to which Countermeasures?
LTD.1	Although the anti-spoofing technology is by this stage well-established, it is nonetheless necessary to conduct regular reviews to ensure that it continues to constitute a reasonable, effective, and acceptable response to a legitimate need. It should be ensured that the system and processes are functioning reliably and as expected, and that performance has not deteriorated over time. The review should ensure that all reasonable steps have been taken, and that all guidelines, industry standards, and data protection authority recommendations have been adhered to. If further possible steps are identified, they should be implemented, and then the review process revisited shortly after.	- End-users	- Human Supervision - Liveness Detection - Multi-modality - Inherently Robust Modalities
LTD.2	If any additions are made to the anti-spoofing technology, they should be closely examined to ensure that they do not cause any privacy or fundamental rights issues. If additions entail capturing more or different data from data subjects, or if they entail processing data in different ways, it should be considered that new impact and context assessments will be required. It may also be necessary to contact the data protection authorities again (this should emerge in the impact assessment procedures).	- End-users	- Human Supervision - Liveness Detection - Multi-modality - Inherently Robust Modalities
LTD.3	A list of “critical incidents” should have been developed (cf. ID.7 above). This should be monitored and added to as necessary (bearing in mind that contexts change over time). As	- End-users	- Human Supervision - Liveness Detection

	ever, should a critical incident occur, , the whole system should be stopped and re-assessed.		- Multi-modality - Inherently Robust Modalities

Table 4. Guidelines for the long-term deployment phase.

5.4. Guidelines on Privacy and Data Protection

This deliverable does not include a template for any form of impact assessment (data protection, privacy, or ethical). Impact assessments should be developed in consultation with relevant local supervisory authorities and standards.[22] Nonetheless, we here provide guidelines on the mitigation of privacy and data protection issues. Any privacy and/or data protection impact assessment would cover at least those issues covered here.

	#	Guideline	Comments Specific to Anti-Spoofing Technologies
Purpose	P&DP.1 **Statement of Purpose**	The purpose of the data processing should be clearly defined and clearly stated.	*The purpose of the biometric system as a whole should be defined and stated (as opposed to the purpose of the anti-spoofing technology in isolation from the rest of system).*
	P&DP.2 **Necessity of Purpose**	The data processing should be genuinely necessary. If it is not necessary (e.g. if there are alternative means of achieving the purpose stated in P&DP.1, which do not require processing personal data) it should not be carried out.	
	P&DP.3 **Purpose Limitation**	Data processing must not extend beyond the stated purpose. Processing for purposes beyond those stated is unlawful (cf. 95/46/EC, arts 6, 10).	
Data Quality	P&DP.4 **Relevance**	Any data collected and processed should be genuinely relevant (i.e. necessary to achieve the purpose of the data processing). If data is to be linked with or accessible to other data sources/systems this should be clearly stated and justified.	*The main factor for anti-spoofing techniques is likely to be minimisation. Those techniques requiring the smallest volume of personal data are to be favoured.*
	P&DP.5 **Accuracy**	Any data collected and processed should be accurate and up to date.	
	P&DP.6	Only the smallest possible amount of data should be collected and processed.	

[22] Privacy and data protection impact assessments are discussed in Annex 1 (§8.4)

	Minimisation	Any redundant data should be destroyed. Personal data should be anonymised wherever possible.	
	P&DP.7 **Rectification**	The data subject has the right that inaccurate data be corrected or removed.	
Storage	P&DP.8 **Security**	Data must be securely stored using any technical, procedural, or organisational means necessary to prevent loss, alteration, disclosure, etc. Centralised databases are to be avoided if possible.	*If data gathered for use in anti-spoofing is not needed for any other stated, legitimate purpose, it need not be retained for longer than it takes to verify that no spoofing is occurring. In many cases it should be possible to erase it immediately after initial use.*
	P&DP.9 **Retention**	Data should be retained only for as long as they are needed.	
Non-discrimination	P&DP.10 **Sensitive Data**	In theory, the following categories of data should not be processed: - Data revealing of racial or ethnic origin; - Data revealing of political opinions; - Data revealing of religious or philosophical beliefs; - Data revealing of trade union membership; - Data concerning health; - Data concerning sex life. However, these categories of data *may* be processed under certain circumstances (defined by 95/46/EC, art 8 (a)-(e)), including if: - The data subject has given explicit consent; - The processing is necessary for protecting the vital interests of the data subject or another person; - The data have been manifestly made public by the data subject; - The processing is necessary for protecting the public interest (e.g. public health).	*In some cases, anti-spoofing techniques derive more or less directly from medical techniques. This suggests a high likelihood of sensitive data being processed. In such cases, explicit consent will be required.*
Fair nes	P&DP.11 **Adequate Justification**	Personal data should only be processed on some good (and legally sound) basis, such as:	*If the addition of anti-spoofing technology to an existing biometric system entails gathering different data, or processing the same data in*

		- The data subject's unambiguous consent; - The processing is necessary for compliance with a legal obligation; - The processing is necessary for protecting the vital interests of the data subject or another person; - The processing is necessary for protecting the public interest (e.g. public health).	*a different way, the data subject should be informed of this, even if informing data subjects may prove difficult.* *Proportionality is always a key consideration. It should be borne in mind that if the addition of anti-spoofing technology to an existing biometric system entails gathering different data, or processing the same data in a different way, the balance of proportionality may shift.*
	P&DP.12 **Transparency**	The processing should be clear and open to the data subject, in order that they can make an informed decision to consent (or not). Prima facie, covert data collection should not take place.	
	P&DP.13 **Access**	The data subject should be able to easily receive details of whether their personal data is held, what data is held, how it is stored and processed, why, and with what justification.	
	P&DP.14 **Right to Withdraw Consent**	Where data are processed on the basis of consent, the data subject retains the right to withdraw consent. There should be clear and simple procedures for signalling the withdrawal of consent.	
	P&DP.15 **Proportionality**	Any processing of personal data should be a proportionate response to a legitimate need. The benefits gained through the processing should offset the intrusion of and risk to the data subject's privacy.	
	P&DP.16 **Automated Decisions**	Data subjects should not be subject to legally (or otherwise) significant decisions taken automatically (i.e. without a competent human making the final decision) on the basis of processing of their personal data.	

The Wider Concept of Privacy	P&DP.17 **Dignity**	People must not be obliged or coerced into situations or activities in which there is a reasonable likelihood of their dignity being compromised. (E.g. procedures should not cause humiliation.[23])	*Being subject to anti-spoofing technology, some people may feel under suspicion or investigation. They may feel such investigation as an affront to their dignity and integrity, as well as an unwarranted intrusion.*
	P&DP.18 **Integrity**	People must not be obliged or coerced into situations or activities in which there is a reasonable likelihood of their physical or psychological integrity being compromised. (E.g. procedures should not be physically invasive or psychological intrusive.[24])	
	P&DP.19 **Autonomy**[25]	People must not be obliged or coerced into situations or activities in which there is a reasonable likelihood of their autonomy being compromised. (E.g. procedures should not undermine the subject's capacity to make free and informed decisions.[26])	

Table 5. Guidelines on privacy and data protection.

[23] This could occur if, for e.g., someone is asked to undress, or reveal a part of the body they deem intimate or sensitive.
[24] This could occur if, for e.g., a sub-dermal procedure is deemed by the subject to compromise their sense of physical unity.
[25] Autonomy is, broadly speaking, the capacity of an individual to make decisions of their own, free from coercion by "external" factors. For discussion in the anti-spoofing context, see D7.2 (§3.2). For more general discussion at a conceptual level, see, e.g., Buss (2008).
[26] This could occur if, for e.g., the alternatives to the use of a certain technology all entail a high cost such as lengthy queuing time, inability to access essential goods or services, and so on.

5.5.Guidelines on Fundamental Rights

Impact assessments should be developed in consultation with relevant local authorities and standards.[27] Nonetheless, in this subsection we provide guidelines on respect for fundamental rights.

	#	Guideline	Comments Specific to Anti-Spoofing Technologies
Respect for the Person	FR.1 **Human Dignity**	Dignity is a fundamental, non-derogable right. It must be respected at all times. No technology or procedure should risk compromising a person's dignity. A person's dignity could be at risk if (*inter alia*) a technology or procedure threatened to objectify them, treat them merely as a means to an end, unduly restricted their autonomy, or shamed or humiliated them in any way.	*It should be considered that anti-spoofing techniques derived from (or at least connected with) medical techniques may be, or be seen to be, a threat to dignity, integrity, and privacy.* *The need for anti-spoofing could be construed as in some way implying guilt or suspicion. Moreover, any technique that might be seen as a threat to integrity should require the data subject's free, informed, and revocable consent.*
	FR.2 **Integrity**	Integrity incorporates both physical and psychological integrity. No technology or procedure should undermine a person's integrity by, for example, being unduly invasive (in a physical or psychological sense). Free and informed consent is an aspect of the right to integrity. If a technology or procedure could plausibly be seen as a threat to integrity, its deployment *must* be subject to free and informed consent.	
	FR3. **Privacy**	The right to privacy is wider than the right to data protection. It includes aspects of dignity and integrity.	

[27] Privacy and data protection impact assessments are discussed in Annex 1 (§8.4).

Non-Discrimination	FR4. **Respect for Diversity**	No technology or procedure should restrict, or make unduly and disproportionately difficult, any person's access to a service or good. Technologies and procedures should be usable by all people regardless of their background and idiosyncrasies.	*Prior to the deployment of any anti-spoofing technology, it should be ensured that it is accessible to all users, regardless of age, gender, background, religion, disabilities, etc. If an anti-spoofing technology cannot be made accessible to all (e.g. gait biometrics cannot be used by wheelchair users), it must be ensured that there are unobtrusive alternatives available, at no disadvantage to the user (i.e. there should not be a significantly greater waiting time, the alternative process should not be more intrusive or time-consuming, the provision of the alternative should not be ostentatious or single out and identify the user as a "special case"). Cultural sensitivities should also be taken into account. If some users will be unhappy to use the technology (e.g. because it requires them to do something that makes them uncomfortable, such as revealing a bodily feature), mitigation strategies or alternatives should be made easily available.* *It may be necessary to introduce testing, prior to deployment, to ensure adequate provision for children, the elderly, or persons with disabilities. Where necessary staff (and possibly users) should receive additional training.*
	FR5. **Rights of the Child**	Children may have particular requirements. Young children may not understand a technology or procedure; they may not be able to give consent on their own behalf; they may, due to their age, be unable to perform the same tasks or meet the same standards as adults. Where children have to use a technology or procedure, provision should be made to ensure that they can, or that viable alternatives are provided.	
	FR6. **Rights of the Elderly**	Older people may have particular requirements. They may be less familiar with a technology or procedure (e.g. ICT); they may, due to their age, be less able to perform certain tasks. Provision should be made to ensure that the elderly nonetheless are able to use a technology or procedure, or else viable alternatives provided.	
	FR7. **Rights of Persons with Disabilities**	Persons with disabilities may have particular requirements. Their specific requirements will vary according to the nature of their disability. Provision should be made to ensure that they are able to use the technology or procedure, or else viable alternatives provided. In the context of biometrics, ISO (2008) offers the following examples of disabilities. 1. **The absence of physical body parts required for the correct operation of a biometric** or its specific instantiation in the system. (Example: missing index finger(s) in an access control system using prescribed fingers.) 2. **The absence of behavioural features required for the correct operation of a biometric** or its specific instantiation in the system.	

		(Example: data subject with no power of speech required to use a voice-activated door entry system.) 3. **Unusable physical body parts required for the correct operation of a biometric** or its specific instantiation in the system. (Example: person with extreme arthritis asked to use a flat plane hand geometry biometric.) 4. **Unusable behavioural features required for the correct operation of a biometric** or its specific instantiation in the system. (Example: data subject in a country with a writing system based on non-Latin alphabet required to use a dynamic signature system designed for Latin alphabets.) 5. An **inability to present the required biometric characteristic in a sufficiently consistent and predictable manner** under the particular conditions of operation. (Example: (i) uncontrollable movement of the eyeball resulting in difficulty in operating an iris recognition system; (ii) person with a speech impediment (e.g. stuttering) asked to use a speaker verification scheme.) 6. An accelerated drift, that is a **change in a characteristic over a period of time in physical or behavioural aspects resulting in increasing difficulty in meeting the matching criteria for an identification or verification**. (Example: data subject with conditions that rapidly age the facial features being verified in certain automatic face verification systems.) 7. An **inability to access, or difficulty with physical access to, the biometric sensor or user terminal**. (Example: data subject or person with a stature not tall enough to access a sensor or user terminal fixed at a specific height.)	

		8. An **inability either to read, due to illiteracy, or to understand the instructions, or to recall the correct procedures, in order to operate the biometric system** successfully. (Example: Forgetting which finger was enrolled in an unattended access control system, and being locked out after three attempts.) 9. **Psychological conditions that prevent the data subject operating the biometric systems correctly**. (Example: Persons with extreme compulsive-obsessive disorder required to use sensors or keypads/keyboards with physical contact.) 10. Conditions, such as those listed above, that result in **disproportionate use of resources**. (Example: Senior citizens who require a longer period of adjustment to changes in context and situation, exceeding the notional time allowed for an authentication.) 11. **Inability to capture biometrics for children or individuals that don't have "standard"-size biometrics**. (Example: Child using a hand geometry reader due to the position or size of the sensor.) No data concerning a person's medical condition should be recorded (e.g. that they *could not use the scanner due to arthritis*).	

Table 6. Guidelines on fundamental rights.

5.6. Contextual Factors Analysis

Prior to the deployment of a technology, an analysis should establish what provisions are needed to account for specific factors of the deployment context. Here we provide guidelines for the identification of ethically salient features in a given context. With the contextual factors analysis completed, it should be easier to determine how best to proceed with the deployment of a biometric/anti-spoofing technology in a given scenario.

#	Contextual Parameters and Guidelines
CF.1 **General Contextual Backdrop**[28]	**Parameter Description**: The *general kind* of context that is at stake. As a rule of thumb, it can be determined by asking someone to respond, more or less instantly, to the question: "what kind of context is this?" Since we are talking at a very general level, the assumption is that a single simple answer will be forthcoming in almost all cases (such as: health; employment; security; border control, etc.). If necessary, the broad use contexts from PD.1 could be used or adapted. The general contextual backdrop will also include the locality and industry, again at very general levels (e.g. Belgium, the petrochemical industry). Establishing the general contextual backdrop will be of assistance in identifying other contextually salient factors (e.g. historical factors, vulnerable groups, etc.). **Examples**: If considering using biometrics to manage access to a school canteen in Stockholm, the general contextual backdrop would be *education in Sweden*. **Guideline**: Establish the general contextual backdrop of the proposed deployment of the anti-spoofing technology.
CF.2 **Stakeholder Interests**[29]	**Parameter Description**: This describes who the main stakeholders are in any given case, and what their interest in that case is. Stakeholders must include *at least*: those *deploying* the anti-spoofing technology; and those *subject to* the anti-spoofing technology. **Examples**: In a border control case, stakeholders include *border authorities* and *passengers*. Border authorities' interests

[28] Derived from Nissenbaum's (2010) parameter "context".
[29] This is an expansion from Nissenbaum's (2010) parameter "actors".

	include *security*; passengers' interests include *convenient, disruption-free, international travel.* **Guideline**: Identify all relevant stakeholders and their specific interests. Having identified relevant groups, if possible they (or their representatives) should be contacted and their views on the proposed deployment solicited. If there are any significant imbalances in power-relations among stakeholders, this should be noted.
CF.3 **Information Types**[30]	**Parameter Description**: This describes what (broadly) the information or data which is to be gathered is *about*, and in what *form* it is processed. **Examples**: Whether personal data is processed; whether it is also sensitive data; whether financial information is gathered; whether biometric information is gathered; etc. **Guideline**: Establish precisely what personal data is to be processed. At this stage, if it is uncertain whether a particular form of data is properly classified as *non-personal*, *personal*, or *sensitive*, always assume the most demanding category (e.g. if unsure whether a certain form of personal data is sensitive or not, assume that it is).
CF.4 **Purpose, Needs, and Musts**	**Parameter Description**: This describes the purpose of the information processing system, and whether there are any necessary or compulsory aspects to the system's deployment. Is the use of biometrics genuinely necessary? Are less intrusive alternatives available? **Examples**: The purpose of the system should be set out clearly, and from this its necessity can be judged. A system managing accessing to a critical infrastructure site, for e.g., will warrant the use of extremely strong identifiers (e.g. biometrics), whereas access to a less critical site (e.g. a library) may not. **Guideline**: The purpose of the anti-spoofing technology should be carefully considered. It should be formally documented, along with a preliminary justification of its necessity. (If the Contextual Factors Analysis is carried out after the Privacy and Data Protection analysis, details from the latter may be used in carrying out the former. However it is better to carry out the Contextual Factors Analysis first.)

[30] Derived from Nissenbaum's (2010) parameter "information types".

CF.5 **Standards and Frameworks**[31]	**Parameter Description**: This describes any established principles constraining or governing the transmission or sharing of information in the "General Contextual Backdrop" (cf. CF.1). These principles, standards, and frameworks may be legal or customary. **Examples**: Legal frameworks for data protection; specific national legislation; industry standards; organisational or voluntary codes of conduct. **Guideline**: Identify all relevant legal frameworks (e.g. regarding data protection), and all relevant industry standards and codes of best-practice.
CF.6 **Vulnerable Groups**	**Parameter Description**: Are any potentially vulnerable groups going to be using the system? A group might be considered vulnerable if, e.g., they have difficulty interacting with a biometric scanner (e.g. some disabled people); they are unable to consent (e.g. minors); they have an unusually high or unavoidable need for the service to which the biometric system manages access (e.g. hospital in-patients); they are likely or liable to feel persecuted or isolated (e.g. ethnic or religious minorities). (All of these points should also be considered in relation to the staff deploying the technology.) **Examples**: Vulnerable groups might include (inter alia): the elderly, minors, the disabled, ethnic or religious minority groups, institutionalised populations (prisoners, in-patients, etc.), people with learning difficulties. **Guideline**: Carefully consider all possible users of the anti-spoofing technology. Will all of these groups be able to easily use it? Will any users be disproportionately disadvantaged by the technology? Will any users be (inadvertently) discouraged, or less likely, to use the technology? If necessary, alternative provision *must* be provided. Will any users of the technology stand in significantly unbalanced power-relations with other stakeholders (e.g. in the workplace)? (All of these points should also be considered in relation to the staff deploying the technology.)
CF.7 **What is at**	**Parameter Description**: To what does the biometric system or anti-spoofing measure manage access? Is it very valuable? In what sense is it valuable (e.g. financial, cultural, security, etc.)?

[31] This is an expansion from Nissenbaum's (2010) parameter "transmission principles".

stake?	**Examples**: An individual's banking details or bank account; an individual's personal files (e.g. on a laptop or smartphone); an organisation's data (e.g. on industrial or personnel matters); a secure location (e.g. critical infrastructure site, border control, private building); a museum or cultural attraction, a religious site; etc. **Guideline**: Compile a comprehensive list of the services, goods, locations, facilities, or opportunities that the biometric system/anti-spoofing technology manages access to. Consider that value may not be solely financial, but could also be, for example, cultural. Is there any security value to the protected services, goods, and opportunities? Consider that what is of little value to one person may be of considerable value to someone else. Having established a comprehensive list, carefully consider who might attempt to gain illicit access, and what the cost of illicit access could be (bearing in mind that "costs" need not be financial).
CF.8 **Cost to Spoofer**	**Parameter Description**: This describes the different costs that a spoofer would incur in developing a spoofing attack. **Examples**: Financial costs; time; effort; special materials; expertise; detailed information; etc. **Guideline**: Building on the analysis from CF.7, carefully consider what the costs and requirements to likely spoofers would be. Think not only in financial terms, but also in terms of time, special materials, technical expertise, training, and so on. Think also in terms of what means the spoofer might attempt to use. Are there "backdoor" vulnerabilities which could be exploited?
CF.9 **Additional Security Features**	**Parameter Description**: This describes any additional security features that could be relevant to the likelihood of there being successful spoofing attacks. (So this includes all security measures not internal to the system itself.) **Examples**: Security measures independent of the anti-spoofing technology, such as: human oversight; deployment of systems within already secure areas; presence of CCTV; presence and proximity of authorities (e.g. police officers); the general level of security in the society (e.g. is there a high threat of, e.g., terrorist attacks or other crime?). **Guideline**: Describe all the additional security measures that are in place in the proposed deployment context. If the proposed anti-spoofing technology is to be used in a deployment context containing other biometric, identification, or security technologies, it's interaction with these and their security features should also be described. It should be considered that, if

	data is to be shared or linked with data from other sources (external databases, e.g.) this may aggravate security concerns.
CF.10 **Cultural Factors**	**Parameter Description**: What cultural norms are relevant to the deployment of the proposed biometric system or anti-spoofing measure? (Note that to some extent this factor may overlap with "Social and Political Factors" and "Historical Factors" (see CF.11 and CF.12).) **Examples**: Norms regarding revealing parts of the body; norms regarding undressing in public; norms regarding touching sensors that others have touched; norms regarding being observed in various ways; norms regarding levels of suspicion (e.g. "why do I have to reveal my identity at all?"); norms regarding anonymity; etc. **Guideline**: List all relevant cultural factors pertaining to the proposed deployment context. In compiling a comprehensive list it will be necessary to consider both the prevailing cultural context as determined by the proposed geographical/civic location, but also all cultural factors stemming from the presence of cultural groups in the community (such as religious groups, immigrant communities, and other minorities).
CF.11 **Societal and Political Factors**	**Parameter Description**: This describes the political situation in the society in which the system is deployed. (This factor may be of interest if the system is deployed by the authorities.) (Note that to some extent this factor may overlap with "Cultural Factors" and "Historical Factors" (see CF.10 and CF.12).) **Examples**: Is the society democratic? Is there societal or political unrest or dissatisfaction in the air? **Guideline**: List all relevant societal and political factors pertaining to the proposed deployment context. These will include the broad political system in place, but also factors such as the public's trust in the authorities, any political or societal tensions, and so on. There may be narratives of privacy or security or the use of biometrics in the media: these should be noted too.
CF.12 **Historical Factors**	**Parameter Description**: This describes any historical factors that could influence the perception of the deployment of the system. (Note that to some extent this factor may overlap with "Cultural Factors" and "Societal and Political Factors" (see CF.10 and CF.11).) **Examples**: Does the end-user have a track-record of responsible use of ID management systems? Are there likely to be any

	longstanding sources of opposition to the deployment (cf. the attitude towards ID cards in the UK). **Guideline**: List all relevant historical factors pertaining to the proposed deployment context. These will include broad historical trends in the society. For example, is there likely to be opposition to the biometric system/anti-spoofing technology due to past problems? Or is the biometric system/anti-spoofing technology likely to be highly acceptable due to previous experience? Historical factors will also include details of the track-record of the end-user proposing the biometric system/anti-spoofing technology; they may also include the history of any relevant stakeholders or vulnerable groups identified in CF.2 and CF.6 respectively (if, for example, historical factors make significant opposition more likely or less likely).

Table 7. Contextual factors analysis.

6. Application to D2.1 Use Cases

In this section we briefly illustrate possible applications of the ethical guidelines. We appeal to the D2.1 use case descriptions, indicating where specific ethical guidelines would be relevant. Given the role of contextually specific factors, we do not go into close detail. That is to say, we propose *that* a guideline should be applied, not *how* it should be applied (i.e. not what specific outcome should emerge).[32]

Each of the following subsections covers a single use case. We quote directly from D2.1 (quotation indicated by indentation), and intersperse (in **boldface** and square brackets "**[**" "**]**") references to relevant guidelines.[33]

6.1. Automated Border Control (D2.1, §2.1)

> The automated gate […] is installed in airport, at port arrival and departure zones, and is equipped with acquisition and processing tools that may vary depending on the selected control process: fingerprint sensor, iris scan cameras, facial acquisition cameras and their flashes, travel document reader (passport, visa), travel ticket reader, integrated PC used for biometric data comparison, single passage detector, etc. **[All PD guidelines are relevant prior to deployment; ID.1; ID.2]**. The combination of detection/acquisition tools are numerous and depend on the security levels required by each airport **[ID.3]**. In this use case, the system has no iris scan cameras, but all the different tools mentioned above are present in the aim to check that the traveller presenting a travel document really is the person to whom this document was issued, and thus verify the person's right to enter or leave the territory **[ID.6; ID.7]**. Moreover, the environment of the system is supervised, with in particular adequate lighting conditions for face acquisition **[ID.4]**.
>
> First, the traveller puts his/her e-passport on the gate reader which reads the biometric data (fingerprints and face) contained in the passport's chip **[ID.5]**. Then the traveller presents successively his/her fingers to the fingerprint sensor and his/her face to the facial camera so that the gate acquires his/her fingerprints and face. A comparison of the traveller's biometric data is performed with those contained in the e-passport's chip. If this comparison leads to a successful match, the gate opens and the passenger is invited to enter the territory. Otherwise, the passenger is denied access and a manual control by a police officer may be carried out **[ID.4; ID.7; once deployed, ID.8; then all LTD guidelines]**.

[32] Ethical issues relating to use cases are discussed in D7.1.

[33] Thus the length of the use case descriptions used here is determined by the details provided in D2.1, and in no way implies that the guidelines are more or less relevant to any particular use case. Rather, we hold that the guidelines are equally applicable to all cases.

6.2. Physical Access Control (D2.1, §2.2)

Physical access control is any mechanism by which a system or a person grants the right to access some area. E.g. a bank employee must normally pass the secured entrance to get to the safe.

The main thing in all security levels, whether it`s a safe in a bank or the copy room of a small company, is that no unauthorized persons could have access to the non-public areas **[All PD guidelines are relevant prior to deployment]**.

The solution for the control of the entrance to a secured area should be automated, i.e. machine controlled, hence personal control by personally knowing the persons, wishing to enter, is not really possible in today's business approaches **[ID.1; ID.2]**. Machine authentication is considered in three levels:

a. What you have: Like Keys, ID – Badges, tokens, etc.
b. What you know: Password, PIN, etc.
c. Who you are: Biometric Authentication.

Most biometric access control systems are using the biometric fingerprint **[ID.6]**. Some are orientated towards biometric face verification. Both of those biometrics and as most other individual biometrics underlie the spoofing and replay attacks and must therefore have personal surveillance to detect eventual spoofing **[ID.3; ID.5; ID.6; ID.7; once deployed, ID.8; then all LTD guidelines]**.[34]

6.3. Logical Access (D2.1, §2.3)

Login computer: is a computer login application **[All PD guidelines are relevant prior to deployment]**. At the login panel, instead of selecting the right profile and entering the associated password, the system automatically recognises the user and logs into the correct account **[Once deployed, ID.8; then all LTD guidelines]**.

Lock screen: locks the computer when the user moves away from it, according to a settable timer which checks for user inactivity (mouse and keyboard) **[All PD guidelines are relevant prior to deployment]**. When the right user comes back and looks at the camera, the system unlocks the computer. The user may also set a periodical timer which checks, regardless of mouse and keyboard inactivity, if the person using the computer is the right user. If not, the system immediately locks the computer **[Once deployed, ID.8; then all LTD guidelines]**.

[34] Nota bene: consideration of ID.4 indicates that a real use case would have to describe where some level of human oversight applies (this is lacking in D2.1).

Internet service: is web-based password manager **[All PD guidelines are relevant prior to deployment]**. When the user goes to a website that requires a login step, the system offers to save the login information. The next time the user wants to access, it automatically logs the user into the website if the right person sits in front of the camera **[Once deployed, ID.8; then all LTD guidelines]**.

The login computer system is based on personal biometric authentication technology working as follow: The user starts up his/her computer and launches the personal biometric authentication software **[ID.1; ID.2]**. He/She then creates a face model which will be saved locally in the computer **[ID.3; ID.5]**. The face model has to be validated by a password (identical to the login password). He/She attributes to the face model the login credentials (user name and password). When the user starts up the computer, he/she will be logged in if the person in front of the camera matches to the registered face model **[ID.6; ID.7]**.[35]

6.4. Mobile Access (D2.1, §2.4)

Instead of typing the PIN number to get access to the smart phone, a biometrical authentication should be used **[All PD guidelines are relevant prior to deployment]**.

Ideally a strong authentication should be used to start the smartphone, but ideally following authentication processes should be running to verify the rightful user at any time, according to the importance (need of strength) of authentication on any application as such **[ID.1; ID.2; ID.3; ID.5; ID.6; ID.7; once deployed, ID.8; then all LTD guidelines]**.[36]

6.5. Covert Identity Verification for Secure Environments (D2.1, §2.5)

There is a need for secure monitoring systems wherein the identity of a subject can be determined not necessarily with their knowledge, and where subjects can be channelled so as to enter through a gateway in which biometric systems could be installed **[All PD guidelines are relevant prior to deployment]**. This could be used in banking or in high security where high value predicates need for extra security (we assume that such uses are consistent with some legitimate basis laid down by law) **[ID.1; ID.2; ID.7]**. […] Naturally, such a high security portal system could use multiple biometrics, especially if identification were to be covert **[ID.3; ID.4; ID.5; ID.6]**. Using gait recognition with such a system could not only mediate acquisition of other biometrics, but also be deployed as a biometric in its own right **[Once**

[35] Nota bene: consideration of ID.4 indicates that a real use case would have to describe where some level of human oversight applies (this is lacking in D2.1).

[36] Nota bene: consideration of ID.4 indicates that a real use case would have to describe where some level of human oversight applies (this is lacking in D2.1).

deployed, ID.8; then all LTD guidelines].

6.6. Medical Confirmation of Subjects before Brain Stimulation Application (D2.1, §2.6)

This use case represents the subject confirmation of a user under a medical brain stimulation application **[All PD guidelines are relevant prior to deployment; ID.1; ID.2]**.

There are four different actors participating in this Use Case:

- Subject
- Supervisor
- Brain stimulation and recording system
- Authentication system

The subject is the person who is receiving the brain stimulation **[ID.5]**. It interacts with the supervisor who is the doctor or scientist in charge of the supervision of the treatment or the experiment **[ID.4]**. This actor may not be present when the brain stimulation system would be easy enough to be handled by the subject himself. That means the stimulation and recording device should be carried, placed and set-up correctly, before the treatment or experiment begins **[ID.3, ID.7]**. The authentication system interacts with the stimulation and recording system by producing authentication results whenever it is requested for.

There is a precondition that shall be met before the use case could finish successfully. The subject has to be previously enrolled in the authentication system. This means that the system has previously stored the biometric signature of the user, so it can be compared with the live data captured during the treatment or experiment, in order to provide an authentication result **[ID.6]**.

The start of the use case is triggered when the subject is available to put on the stimulation and recording electrodes. Once it happens the basic flow of action for this use case is the following:

- The doctor or scientist assists the subject in order to place the brain stimulation electrodes. These are the same ones that allow the monitoring of EEG signals.
- The stimulation system proceeds to authenticate the subject according to the received EEG signal and the stored biometric signature.
- If the authentication leads to a positive result, then the system starts the brain stimulation session during the pre-established period of time, unless the doctor or scientist decides to interrupt it.

- The system might carry out additional checks during the brain stimulation session, according to the configuration of the system.

An alternate flow of actions may happen if the authentication of the subject under brain stimulation does not lead to a positive result, then the stimulation device would block thus avoiding that a subject not enrolled in the system could use the brain stimulation system **[Once deployed, ID.8; then all LTD guidelines]**.

6.7. Trusted Telepresence Brain-Computer Interface (BCI) Application (D2.1, §2.7)

[All PD guidelines are relevant prior to deployment; ID.1; ID.2] The main actor in this use case is the user who requests the access the Telepresence BCI system. This request triggers the beginning of the use case. He/She interacts with a BCI system through an electrophysiological sensor that he/she is wearing. The BCI system collects the electrophysiological signals that are translated to commands that are sent through the network to the remote place **[ID.3]**. In the remote place there is the system in charge of controlling the physical avatar, collecting the information sent by the sensors, and driving the actuators that may be installed there.

The BCI system incorporates an authentication module that uses the electrophysiological biometric signals for authenticating the user who is accessing the system **[ID.6]**. It provides an authentication result that allows the Telepresence BCI application to handle the user registration in the system **[ID.5]**. […]

There is a precondition that shall be met before the Use Case can be accomplished successfully. The Telepresence BCI user has to be previously enrolled in the authentication system, so a biometric signature constructed by features of the user's electrophysiological signal has to be available for the authentication module **[ID.4]**. Then this signature is confronted with the live data received by the BCI system in order to output the final authentication result.

The basic flow of action for this use case is as it follows:

- The user wears the ECG and EEG sensor.
- The system waits for the ECG and EEG signals to be received and calibrated correctly.
- The system carries out the authentication process according with the received signals and the stored signature of the claimed user.
- If the authentication process leads to a positive result, the user can access to all the services of Telepresence BCI application.
- The authentication process of the electrophysiological biometric signals is repeated periodically according to the configuration of the system.

Two main circumstances that can modify the Use Case basic flow are contemplated. On one hand, a problem might occur in the reception or calibration of the ECG and EEG **[ID.7]**. On the other hand, the authentication process might lead to a negative result.

In both cases the identity of the subject that is using the system cannot be assured. Using a restrictive policy, the system would not allow the subject to use the system until the authentication process completes successfully **[ID.5]**. If a laxer policy is applied the system might allow the subject to use certain functionalities of the Telepresence BCI system, but assigning to the avatar a non-authenticated state. In this case, people who interact with the avatar could be aware of the abnormal situation with this user **[Once deployed, ID.8; then all LTD guidelines]**.

6.8. Border Control (D2.1, §3.1)

[All PD guidelines are relevant prior to deployment; ID.1; ID.2] At a border control, a manual control booth is primarily designed to check biometric passports and visas. The tools of the booth are used to compare document data with a national or European database. Here, the biometric tool is a fingerprint sensor (this sensor may be dedicated to single fingerprint or multiple fingerprints). It will aim to acquire and compare the fingerprint of the passenger with the one contained in its biometric passport.

The system consists then of a control booth with an operator in it, a control screen, a biometric passport reader and a fingerprint sensor **[ID.3; ID.5; ID.7]**. The environment of the system is supervised **[ID.4]**. Communications between components are secure, so that data protection is ensured (for instance with a digital signature for data integrity check or encrypted data exchange using secure tunnelling) **[ID.6; once deployed, ID.8; then all LTD guidelines]**.

6.9. Delivery of Pharmaceuticals (D2.1, §3.2)

This use case represents the delivery of pharmaceuticals for a patient using a biometric healthcare card. A terminal allowing Match-on-Card is used at the pharmacy to prove the patient's rights **[All PD guidelines are relevant prior to deployment; ID.1; ID.2]**.

All the beneficiaries of the social protection are previously enrolled so that their biometrics can be stored on a personal ID card dedicated to benefit from health care, social security reimbursements, pharmaceuticals delivery, etc. **[ID.3; ID.4; ID.5; ID.6]**. Here is considered the latest of these use cases, i.e. pharmaceuticals delivery.

Once enrolled by collecting all the necessary alphanumeric data (surname, first name, date of birth, gender…) and biometrical data (fingerprint, photo), a beneficiary is given a personalised card with his/her photo printed on it **[ID.4; ID.5]**. All these data are typically digitalised in the

microprocessor of the card, providing encryption and data protection. […] Identity control is ensured by an ID data verification system which basically use fingerprint sensors for checking the card ownership validity **[ID.6]**.

When presenting to a pharmacy for claiming medicines previously ordered by a doctor, a patient is required to present his/her electronic medical card. This card is introduced in a reader equipped with a fingerprint sensor, on which the patient is invited to put his/her index **[ID.5]**. Biometric data contained in the chip of the card are then compared to the fingerprint acquired in live in order for the patient to prove his/her identity and claim his/her medical rights. By performing the matching and certifying the genuineness of the data, the card is self-sufficient and there is no need of using a centralised biometric database **[ID.3; ID.7]**.

According to the pharmacy configuration, the terminal reading the card may be mobile or landline, provided by a pharmacist or not. It can even be located outside the pharmacy in the case of an open-24h-a-day pharmacy with a take-away service in which the delivery could be automatic or assured by a night-working druggist. So this terminal is definitely placed in an unsupervised environment with hard acquisition conditions (as regards light, rain, humidity, temperature…) **[ID.5; once deployed, ID.8; then all LTD guidelines]**.

6.10. Banking System (D2.1, §3.3)

The system consists in an ATM equipped with a contactless sensor, specifically a fingerprint-vein sensor, which means that even if vein is the main biometrics, fingerprint could be used for multi-modal fusion for instance **[All PD guidelines are relevant prior to deployment]**. Otherwise, all aspects of the ATM are those of a standard one: interactive screen, credit card reader, numeric pad… The system is also connected to a local database for all ATMs of the bank using this technology (the database could be extended to all ATMs on a national scale) **[ID.1; ID.2; ID.3]**.

For a cash withdrawal, the end-customer has to put his/her credit card in the reader first. Then, instead of entering a pin code, the user performs the verification of card's ownership by moving his/her hand in the sensor **[ID.5]**. The sensor scans the biometrics of the client, processes the image and turns it into a template thanks to coding algorithms. All these data are mapped and encrypted so that the information is protected. A back office system provides from the database the registered data, by using the card signature for instance, that may be next compared to this template to make a one to one check at the ATM. If the matching is successful, the cash withdrawal is authorised by the ATM, otherwise the first of the three classical attempts is blown **[ID.6]**.

On the client side, the use of vein recognition as a means of authentication is a benefit for the end-customer since the transaction is accelerated (showing his/her hand being faster than

entering a code pin). Moreover, vein network ensures the privacy of the user, so that the client is more willing to use this technology **[ID.7; once deployed, ID.8; then all LTD guidelines]**.[37]

[37] Nota bene: consideration of ID.4 indicates that a real use case would have to describe where some level of human oversight applies (this is lacking in D2.1).

7. Conclusion

Guidelines should be sufficiently broad as to apply in many situations, yet flexible enough that they can be of practical use in very different contexts. The impact of contextual factors in determining ethical issues has been the chief finding of WP7 research. The guidelines we have presented in this document reflect this.

In §3 we presented summaries of the applicable European policy and fundamental rights framework. In §4 we presented a summary of the ethical issues researched in WP7. These summary sections are supported by more detailed research, which is included as annexes to this document (§8, §9).

We have presented guidelines to be followed at the *pre-*, *initial*, and *long-term deployment phases* (§§5.1-5.3). The justification for this approach is presented in §2, which also indicates the purpose of the guidelines, the groups they target, and the advantages of following them.

The deployment phase guidelines are supplemented with guidelines on *privacy and data protection* (§5.4) and *fundamental rights* (5.5). These are, to a large extent, derived from the research presented in §5 (also §8), §6 (also §9) respectively. Finally, the deployment phase guidelines are also supplemented by a *contextual factors analysis* (§5.6), which supports the identification, and response to, significant contextual factors. In §6 we have indicated how the guidelines could be used in the Tabula Rasa use cases.

References

Tabula Rasa Deliverables

D2.1 *Definition of Use Cases*
S. Caillebotte, S. Revelin (MORPHO), J. Acedo (STARLAB), G. Florey (KEYLEMON), J. Bustard, M. Nixon (USOU), M. Dambrauskaite (BIOMETRY), 2011

D7.1 *Environmental Scanning Report*
H. Ashton, E. Mordini, A.P. Rebera (CSSC), 2011.

D7.2. *Interviews Report*
E. Mordini, A.P. Rebera (CSSC), 2012.

D7.3. *Ethical Monitoring Report (Intermediary)*
E. Mordini, A.P. Rebera (CSSC), 2012.

D7.4. *Workshop Report*
E. Mordini, A.P. Rebera (CSSC), 2012.

Policy Documents, Standards and Guidelines, Grey Literature (etc.)

ISO, 2008 *Information technology – Biometrics – Jurisdictional and societal considerations for commercial applications – Part 1: General guidance*. International Organization for Standardization, 2008 (ISO/IEC TR 24714-1:2008).

A29WP, 2011 *Opinion 15/2011 on the definition of consent*. Article 29 Data Protection Working Party, 2011 (01197/11/EN, WP187).

A29WP, 2012a *Opinion 2/2012 on facial recognition in online and mobile services.* Article 29 Data Protection Working Party, 2012 (00727/12/EN, WP192).

A29WP, 2012b *Opinion 3/2012 on developments in biometric technologies.* Article 29 Data Protection Working Party, 2012 (00720/12/EN, WP193).

A29WP, 2102c *Opinion 01/2012 on the data protection reform proposals*. Article 29 Data Protection Working Party, 2012 (00530/12/EN, WP191).

CoE, 1997 *Recommendation No. R (97) 5 of the Committee of Ministers to Member States on*

the Protection of Medical Data, Council of Europe Committee of Ministers, R97(5).

CoE, 2011 *The need for a global consideration of the human rights implications of biometrics*. Council of Europe. Parliamentary Assembly Resolution 1797, 2011.

ECHR *Convention for the Protection of Human Rights and Fundamental Freedoms*. Council of Europe, 1950.

EU, 2007 *Explanations Relating to the Charter Of Fundamental Rights*. Official Journal of the European Communities, 2007/C 303/02.

F'sight, 2013 *Future Identities – Changing identities in the UK: the next 10 years*, Foresight Future Identities (2013), Final Project Report, The Government Office for Science, London.

CFREU *Charter of Fundamental Rights of the European Union*, Official Journal of the European Communities (2000/C 364/01).

95/46/EC *Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data*. Official Journal of the European Communities (L 281/31).

EC, 2010 *A comprehensive approach on personal data protection in the European Union*, Communication from the Commission to the European Parliament, the Council, the Economic and Social Committee and the Committee of the Regions, COM(2010) 609.

EC, 2012 *Proposal for a Regulation of the European Parliament and of the Council on the protection of individuals with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation)*, COM(2012) 11.

Books, Articles (etc.)

Agamben, G. (2008). No to bio-political tattooing. *Communication and Critical/Cultural Studies*, 5(2), 201-202. Reproduced from *Le Monde* (2004, January 10).

Allen, A. (1988). *Uneasy Access: Privacy for Women in a Free Society* (Totowa: Rowman and Littlefield).

Bellamy, R. (2012). “The Inevitability of a Democratic Deficit”, *Key Controversies in European Integration*, Hubert Zimmermann, Andreas Dür, eds. (London: Palgrave), 64-71.

Bok, S. (1989). *Secrets: On the Ethics of Concealment and Revelation* (New York: Vintage Books).

Bonde, J-P. (2011). “The European Union’ Democratic Deficit: How to Fix It”, *Brown Journal of World Affairs*, 17(2).

Buss, S. (2008). “Personal Autonomy”, *The Stanford Encyclopedia of Philosophy* (Winter 2012 Edition), Edward N. Zalta (ed.), at: http://plato.stanford.edu/entries/personal-autonomy/.

Bynum, T.W. (1997). “Anonymity on the Internet and Ethical Accountability”, *The Research Center on Computing and Society*, at: http://www.southernct.edu/organizations/rccs/oldsite/resources/research/global_info/bynum_anonymity.html.

De Hert, P. (2005). “Biometrics: legal issues and implications”, at: http://www.statewatch.org/news/2005/apr/jrc-biometrics-paul-de-hert.pdf.

De Hert, P. (2012). “A Human Rights Perspective on Privacy and Data Protection Impact Assessments”, in Wright and De Hert (2012), 33-76.

Etzioni, A. (1999). *The Limits of Privacy* (New York: Basic Books).

Gavison, R. (1980), “Privacy and the Limits of Law”, *Yale Law Journal*, 89, 421-71.

Grijpink, J. (2001). “Privacy Law: Biometrics and Privacy”, *Computer Law and Security Report*, 17(3), 154-160.

Hume, D. (1960[1748]). “On the Original Contract”. In: *Social Contract: Locke, Rousseau, Hume* (New York: Oxford University Press).

Joyner, C.C., Lotrionte, C. (2001). “Information Warfare as International Coercion: Elements of a Legal Framework”, *European Journal of International Law*, 12, 825-865.

Kant. I. (1997[1785]). *Groundwork of the Metaphysics of Morals*, M. Gregor (ed.) (Cambridge: Cambridge University Press).

Kindt, E. (2007). “Biometric applications and the data protection legislation”, *Datenschutz un*

Datensicherheit 31(3), 166-170.

Largent, E., Grady, C., Miller, F.G., Wertheimer, A. (2012). "Misconceptions about Coercion and Undue Influence: Reflections on the Views of IRB Members", *Bioethics*, DOI: 10.1111/j.1467-8519.2012.01972.x

Liu, Y. (2008). "Identifying Legal Concerns in the Biometric Context", *Journal of International Commercial Law and Technology*, 3(1), 45-54.

Locke, J. (1975[1689]). *An Essay Concerning Human Understanding*, P.H. Nidditch (ed.) (Oxford: Oxford University Press).

Marx, G. (1999). "What's in a Name? Some Reflections on the Sociology of Anonymity", *The Information Society*, 15(2).

Maschke, K.J. (2003). "Proxy Research Consent and the Decisionally Impaired: Science, the Common Good, and Bodily Integrity", *Journal of Disability Policy Studies*, 13(4): 254-259.

Mordini, E. (2011). "Policy Brief on: Whole Body-Imaging at Airport Checkpoints: the Ethical and Policy Context", in von Schomberg, R. (ed.), *Towards Responsible Research and Innovation in the Information and Communication Technologies and Security Technologies Fields* (Luxembourg: Publications Office of the European Union), 165-209.

Mordini, E. and Ashton, H. (2012). "The Transparent Body: Medical Information, Physical Privacy and Respect for Body Integrity", in Mordini, E. and Tzovaras, D. (eds), *Second Generation Biometrics: The Ethical, Legal and Social Context* (Dordrecht: Springer, 2012), 257-283.

Mordini, E., and Massari, S. (2008). "Body, Biometrics and Identity", *Bioethics*, 22, 488-498.

Mordini, E. and Rebera, A.P. (2012). "Conceptualizing the Biometric Body", in Bus, J. et al (eds), *Digital Enlightenment Yearbook 2012* (Amsterdam: IOS Press), pp. 265-273.

Motha, S. (2012). "The Debt Crisis as Crisis of Democracy", *Law, Culture and the Humanities*, 8(3), 390-397.

Nicolaïdis, K. (2013). "European Demoicracy and Its Crisis", *Journal of Common Market Studies*, 51(2), 351-369.

Nissenbaum, H. (1999). "The Meaning of Anonymity in an Information Age", *The Information*

Society, 15, 141-144.

Nissenbaum, H. (2010). *Privacy in Context: Technology, Policy, and the Integrity of Social Life* (Stanford, CA: Stanford University Press).

Rebera, A.P., Bonfanti, M.E. and Venier, S. (2013). "Societal and Ethical Implications of Anti-Spoofing Technologies in Biometrics", *Science and Engineering Ethics*, at: http://link.springer.com/article/10.1007/s11948-013-9440-9.

Regan, P.M. (1995). *Legislating Privacy: Technology, Social Values, and Public Policy* (Chapel Hill: The University of North Carolina Press).

Solove, D. (2008). *Understanding Privacy* (Cambridge, MA: Harvard University Press).

Van der Ploeg, I. (1999). "The illegal body: 'Eurodac' and the politics of biometric identification", *Ethics and Information Technology*, 1(4), 295-302.

Van der Ploeg, I. (2002). "Biometrics and the body as information: normative issues in the sociotechnical coding of the body", in: D. Lyon (ed.) *Surveillance as Social Sorting: Privacy, Risk, and*
Automated Discrimination (New York: Routledge), 57-73.

Van der Ploeg, I. (2007). "Genetics, biometrics and the informatization of the body", *Ann Ist Super Sanità* 43(1), 44-50.

Westin, A. (1967). *Privacy and Freedom* (New York: Atheneum).

Wright, D. and De Hert, P. (2012). "Introduction to Privacy Impact Assessment", in Wright and De Hert (2012), 3-32.

Wright, D. and De Hert, P. (eds). (2012). *Privacy Impact Assessment* (Dordrecht: Springer).

Wright, D. and Mordini, E. (2012). "Privacy and Ethical Impact Assessment", in Wright and De Hert (2012), 397-418.

8. Annex 1 - European Policy and Fundamental Rights Framework

Deliverable D7.1 (§5) offers an inventory of regulations, reports, working papers, academic articles (etc.) relating to the ethical, legal, and policy implications of spoofing prevention. D7.1 remarks:

> [T]here is very little (if any) existing policy on countermeasures for spoofing and the majority of the documents we mention are those relating to biometrics, security and human rights in general. (23)

This remains the case today: there is no policy specific to spoofing.

The European legal framework concerning biometrics (and so spoofing) is firmly based around data protection and fundamental rights legislation, in particular the *European Data Protection Directive* (95/46/EC), the *European Convention on Human Rights* (ECHR), and the *Charter of Fundamental Rights of the European Union* (CFREU). In this section we revisit the provision of these documents (§8.1), address important issues regarding their interpretation (§8.2), and consider possible implications of the impending revision of European data protection legislation (§8.3).

8.1. Data Protection and Fundamental Rights in Europe

Although not all anti-spoofing provision involves the processing of personal data (e.g. human supervision of scanners), many anti-spoofing technologies do.[38] As such, the ECHR (art. 8) and the CFREU (art. 8), which set out the rights to *respect for private and family life* and *protection of personal data* (respectively) come into play. The former represents a general commitment to privacy, but has been interpreted in case law as covering personal data. The latter is more explicit.

1. Everyone has the right to the protection of personal data concerning him or her.
2. Such data must be processed fairly for specified purposes and on the basis of the consent of the person concerned or some other legitimate basis laid down by law. Everyone has the right of access to data which has been collected concerning him or her, and the right to have it rectified.
3. Compliance with these rules shall be subject to control by an independent authority.

(CFREU: art. 8)

[38] Or, at least, they process data which, combined with other similar data, could identify an individual.

The above right is presently protected by the Data Protection Directive (95/46/EC). Table 8 summarises the major data protection principles enshrined in 95/46/EC.[39]

Principle	Description	Provision in 95/46/EC
Necessity	Processing of biometric data (or other personal data gathered for anti-spoofing) should be a necessary means to a clearly stated, legitimate end. If, therefore, a less intrusive alternative to is available, it is to be preferred.	Article 6
Proportionality	Processing should be a proportionate response to the clearly stated, legitimate need. The value of the processing must offset the intrusion/risk to the data subject's rights.	Article 6
Purpose Limitation	Processing should be carried out only for a clearly stated, legitimate end. Any other processing is prima facie unlawful.	Articles 6, 10
Fairness	Processing should be transparent to data subjects, leaving them the right to object to it, withdraw consent, know how their data is stored, handled, shared, etc. Prima facie, covert/ deceptive data gathering is unlikely to be fair.	Articles 6, 10, 12, 14, 15
Data Minimisation	As little personal data as possible should be processed and retained, consistent with the legitimate purpose of the processing. Excess or redundant data should be destroyed.	Article 6
Accuracy and Relevance	Any personal data stored/processed should be accurate and relevant. The subject has the right to have their data corrected or removed if it is inaccurate.	Articles 6, 12, 14
Data Storage	Data should be stored securely using any technical, procedural, or organisational means necessary to protect against destruction, loss, alteration, disclosure, etc.	Articles 6, 12, 16, 17
Retention Period	Personal data should be stored for no longer than necessary.	Articles 6
Justification	Personal data should only be processed if (inter alia):	Articles 7, 13
	(a) The data subject has given "unambiguous" consent;	
	(b) It is necessary for compliance with a legal obligation;	
	(c) It is necessary to protect the vital interests (e.g. health) of the data subject;	
	(d) It is in the public interest.	
Special Categories of Data/Non-discrimination	Processing personal data which could ground unlawful discrimination (e.g. on the basis of race, health, etc.) is prohibited. The prohibition is lifted if (inter alia):	Article 8
	(a) The data subject has given "explicit" consent;	
	(b) Processing is necessary to protect the vital interests (e.g. health) of the data subject or another person;	
	Processing relates to data which the data subject has "manifestly made public".	

[39] This is only a summary. For fully accurate details consult 95/46/EC itself.

Transparency	Processing should be open and clear so that the data subject can make an informed assessment of the processing.	Articles 10, 12, 14
Automated Decisions	No data subject should be subject to a decision based solely on automated processing of their personal data, if that decision will have a legal or otherwise significant effect.	Article 15

Table 8. Data protection principles.

Beyond privacy and data protection, the following fundamental rights are relevant to biometrics and spoofing.[40]

Fundamental Right	Connection with Spoofing	Provision
Human Dignity	Anti-spoofing procedures must respect human dignity (e.g. a means to gather data cannot involve humiliation, e.g. by revealing intimate parts of the body in public).	CFREU, Article 1
Integrity of the Person	Anti-spoofing procedures must respect physical and psychological integrity (e.g. must not be invasive).	CFREU, Article 3
Prohibition of inhuman or degrading treatment	As per *Human Dignity* and *Integrity of the Person*.	CFREU, Article 4 ECHR, Article 3
Privacy/Private Life	Privacy goes beyond data protection, involving concepts such as dignity, integrity, and autonomy. Anti-spoofing procedures should not compromise these values.	CFREU, Article 8 ECHR, Article 8
Non-discrimination	Anti-spoofing procedures should not involve or lead to illegitimate discrimination (e.g. FAR/FRR rates should not disproportionately disadvantage any particular group).	CFREU, Article 21 ECHR, Article 14
Respect for Cultural, Religious, Linguistic Diversity	Anti-spoofing measures should be respectful of diversity (e.g. voice recognition technologies should be able to cope with regional variations of accent).	CFREU, Article 22

[40] Note that we refer only to the two main European fundamental rights instruments, the CFREU and ECHR. For further details on the grounding and interpretation of the rights set out in the CFREU, see the *Explanations Relating to the Charter of Fundamental Rights* (Official Journal of the European Union, 2007/C 303/02).

Rights of the Child	Technological, procedural, or organisational anti-spoofing provision should be mindful of the possibility that children may be subject to the system, and of any additional fundamental rights concerns this may provoke (e.g. legal status of minors re consent).	CFREU, Article 24
Rights of the Elderly	Technological, procedural, or organisational anti-spoofing provision should be mindful of the possibility that older people may be subject to the system, and of any additional fundamental rights concerns this may provoke (e.g. less well-defined biometric features).	CFREU, Article 25
Integration of Persons with Disabilities	Technological, procedural, or organisational anti-spoofing provision should be mindful of the possibility that persons with disabilities may be subject to the system, and of any additional fundamental rights concerns this may provoke (e.g. difficulty in engaging with a sensor).	CFREU, Article 26

Table 9. Main fundamental rights relevant to anti-spoofing

8.2. Interpretation of the Framework

This section discusses some of the more contentious points arising above. These concern:

- The *nature* of biometric data;
- The concept of *proportionality*;
- The concept of *consent*.

8.2.1. Is Biometric Data Personal Data? Is Biometric Data Sensitive Data?

95/46/EC governs the processing of personal data. “Personal data” and “processing” are defined in the European framework, thus:

> For the purposes of this Directive […] “personal data” shall mean any information relating to an identified or identifiable natural person (“data subject”); an identifiable person is one

> who can be identified, directly or indirectly, in particular by reference to an identification number or to one or more factors specific to his physical, physiological, mental, economic, cultural or social identity. (95/46/EC, Art. 2)
>
> "[P]rocessing of personal data" ("processing") shall mean any operation or set of operations which is performed upon personal data, whether or not by automatic means, such as collection, recording, organization, storage, adaptation or alteration, retrieval, consultation, use, disclosure by transmission, dissemination or otherwise making available, alignment or combination, blocking, erasure or destruction. (95/46/EC, Art. 2)

In D7.3 we discussed the difference between "personal" and "sensitive" data. Sensitive data is defined as a special sub-category of personal data. Sensitive data is:

> [P]ersonal data revealing [of] racial or ethnic origin, political opinions, religious or philosophical beliefs, trade-union membership, and […] data concerning health or sex life. (95/46/EC, Art. 8)

In D7.3 we concluded that biometric data is *always* personal data, and *sometime* sensitive. In this section we add to the question of when biometric/anti-spoofing data is personal and/or sensitive.[41] The issue is important because, from both a legal and an ethical point of view, anti-spoofing technologies are likely to be judged acceptable or not based on a calculation of proportionality; and proportionality is directly affected by the nature of the data (given, e.g., the greater risks from processing sensitive data).

Let us first settle the question of whether biometric data is *personal*. According to, for example, Grijpink (2001), it is necessary to distinguish "person-related detail" from "personal detail". Person-related detail is, in some way, derived from a person's body. As long as some detail or datum stems from a person's body, it may be considered *person-related*. Personal detail, on the other hand, is detail that can actually be *traced back* to an individual. So the condition on being "personal" detail is stricter than simply being "person-related": a datum goes beyond being merely person-related if, from it, it is possible to work back to its source (i.e. the person to whom it is related). Biometric data is by definition person-related, but may or may not be personal detail.[42]

From a legal point of view, Grijpink's distinction between "person-related" and "personal"

[41] We use "biometric/anti-spoofing data" as a convenient shorthand. We do not mean to imply that data gathered for anti-spoofing is always biometric.

[42] "An anonymous biometric characteristic, i.e. a detached biometric template without anything in common with the source, cannot be regarded as a personal detail because it cannot be traced back to the person from whom the measured value originated, or this can only be done with disproportionate effort" (Grijpink, 2001: §4).

information is arguably irrelevant.[43] At root however, Grijpink highlights a key point: the importance we attach to information deriving from an individual should be sensitive to the possibility of working back *from* that information *to* the individual.

This important point is reflected in 95/46/EC, since surely the rationale of employing the concept of an "identifiable person" is precisely to distinguish information that can be traced to an individual from information which, whether or not it *derives* from a specific individual, cannot. The problem, which appears to be inherent in the wording of the directive, is that the modal scope of "identifiable" in the phrase "identifiable natural person" is not specified. Are we talking about a person who is identifiable under any circumstances whatsoever? Or only a person who is identifiable to others who have the appropriate equipment and background knowledge? Are we talking about what is possible right now? Or do we also include what may be possible in the future?

Now this may seem a somewhat arcane point, but it is extremely important. Grijpink appeals to the notion of "disproportionate effort" (meaning the effort to work back from a template to an individual). But this is problematic. Disproportionate to what? If a piece of information is somehow crucial to accessing a large amount of money, then very great efforts to trace it back to the individual from whom it derives might not be disproportionate, given the amount of money at stake. And it should also be considered that what is extremely difficult today may, in one year's time, be relatively easy (e.g. if technology progresses sufficiently).

To settle this point would require developing a far more robust account of "personal data" and the concept of an "identifiable person", supplemented with close analysis of relevant legal and judicial decisions. However we can safely leave the legal question open. For even if legal analysis were to reveal that, say, anonymised biometric templates are not personal data, it would still be true that there is a reasonable likelihood that in the future it will be possible to break that anonymisation. Therefore there is an ethical imperative to handle such data with caution, independently of whether the letter of data protection law demands such treatment. (Consider also that anonymisation is itself an act of data processing, hence 95/46/EC applies at the point of anonymisation.) Here the ethical imperative is potentially more demanding than the legal imperative. For the purposes of these guidelines, we have undertaken to always adopt the stricter of two options. Hence, biometric data is here always considered as personal data, and we recommend processing it accordingly. (Even if anti-spoofing data is not personal, it will likely be processed alongside personal data (i.e. the biometric data), and hence data protection law applies to the system as a whole.)

[43] 95/46/EC defines "personal data" as "relating to an identified or identifiable natural person". Does biometric data relate to an identifiable natural person? Yes. What about biometric templates? Again, given that the template is related to an identifiable natural person—after all, the very *purpose* of the template is to make possible identification or verification—it must be considered as personal data (cf. De Hert 2005; Liu 2008). (It is also worth noting that the ISO guidelines consider biometric data as personal (ISO, 2008: v).)

The question of whether biometric data is *sensitive* turns, for the most part, on the issue of whether it can be reasonably considered as revealing of either health, race or ethnicity.[44] The most likely cases are facial images revealing skin colour, or varieties of biometric data potentially revealing of health conditions (on the latter see: Mordini & Ashton, 2012).

The debate as to whether biometric data reveals health conditions is far from settled. The debate is, at root, an issue of the distinction between theoretical possibility and actual occurrence. In D7.2, one interviewee, Alexander Nouak, claimed that medical data could be discerned from raw data from iris scanning. However another interviewee, Jim Wayman, claimed not have seen any reliable research proving such claims.

The argument is on-going; but a couple of important points provide a foundation for the guidelines.

(1) Firstly, some kinds of raw biometric data are almost certainly sensitive. For example, facial images often reveal racial background (or at least whether one is white, black, etc.). To the extent that such data is revealing in that way, it is sensitive.

If course, not all images will constitute sensitive data. Blurred and distant facial images are unlikely to constitute personal data because they do not allow the recognition of the individuals concerned (A29, 2012a: 4). *A fortiori*, such images are not likely to be sensitive, sensitive data being a subcategory of personal data. In this relation we should consider different kinds of anti-spoofing data. For example, various forms of "challenge-response" liveness-detection require the data subject to perform an action: e.g. roll a finger across a scanner (fingerprint); blink (iris); or read a specific phrase out loud (voice). Neither the performance of these actions nor the data gathered therefrom, need necessarily generate personal data. Sensitive data is set apart for special treatment in the legislation in recognition of the fact that it can be a source of illegitimate discrimination (95/46/EC, preamble 33). Now clearly, the inability to perform the challenge-response actions just described could be a source of discrimination. So whether or not data of this kind is personal, it *can* be a source of illegitimate discrimination on the grounds of health condition. Therefore, the data should from an ethical perspective, be treated very carefully. (Again, even if the strict legal requirement falls short of demanding that such data be handled as if it were sensitive, the ethical imperative outweighs and overrules it—at least for the purposes of these guidelines.)

If the above reasoning is sound, there is a gap in the legislation.[45] The lesson to be drawn, for

[44] Other potential sources of illegitimate discrimination (e.g. religious beliefs) could be revealed by, say, an image of an individual (e.g. if it shows religious paraphernalia such as a crucifix).
[45] A gap which is not addressed in the proposed Data Protection Regulation (EC, 2012).

present purposes, is that the guidelines should promote compliance with the *spirit* of the data protection principles to ensure that, where the wording leaves room for doubt as to the interpretation – and it has to be said that the definition of "personal data", for example, is not clear[46] – high standards are nonetheless maintained.

(2) The second point concerns health-related data and again stems from a lack of clear terminology.

In D7.2. Jim Wayman offered the argument that, while one may be able to demonstrate a *correlation* between a certain sort of fingerprint and a certain sort of genetic disease, a mere statistical correlation does not show, for any give person with the relevant sort of fingerprint, that they have the genetic disease. Therefore, the fingerprint does not indicate the presence of the disease. Now as a matter of logic, this is unobjectionable. However, as matter of data protection, it is open to objection.

The objection focuses on the notion of "health data" (or "medical data" or "data revealing of health", etc.). Surely, one might suppose, the concept of "health data" is not so narrow as to only include data which could be used as a basis for a *reliable diagnosis* of a health condition. For example, there are in the UK (and presumably elsewhere) posters showing the silhouette of a person from a side-on perspective. These posters inform people that if their waist is as big or bigger than the one depicted, they should consider being tested for diabetes. Now the fact that one has a waist of a certain size obviously does not entail having diabetes. But the rationale behind the poster campaign is that there is a sufficiently robust correlation between waist-size and diabetes that the number of diabetics that will be successfully diagnosed and as a result of the poster campaign is worth the time and effort of checking the many people whom come forward for testing having the relevant waist-size but no diabetes. Now, bearing this in mind, is waist-size health data? In at least *some* circumstances – for example when is used for purposes of pre-screening for a health condition – it surely is.[47]

95/46/EC gives no definition of "health data": however the proposal for a new data protection Regulation does.

[46] It is unclear because it depends upon the understanding one attaches to the phrase "identifiable person". "Identifiable" introduces the matter of *possibility* (an issue we discussed above). So does 95/46/EC apply to data from which an individual could be *easily* identifiable, *relatively easily* identifiable, *possibly* identifiable? Not clear. Moreover, as we stressed above, possibilities change (in the sense that what is possible now may not have been possible 15 years ago; and what is impossible now may be possible in the future). How can we make reasonable judgements about this? The matter is extremely complicated. Note that the talk given by Anders Sandberg at the Tabula Rasa Workshop in Rome (May 2012) addressed the related question of how we should think about the privacy rights of people from the past, and how we should think about our own privacy in relation to future generations. See D7.4.

[47] Further related arguments concern the linkability of databases and data-mining. It is true that a single set of, say, dactyloscopic data, could be harmless; yet to the extent that it is possible to link this data with other sets of data, very significant information about a given subject could be derived.

> "data concerning health" means any information which relates to the physical or mental health of an individual, or to the provision of health services to the individual. (EC, 2012: 42)

And the Council of Europe's Recommendation on the protection of medical data states that:

> the expression "medical data" refers to all personal data concerning the health of an individual. It refers also to data which have a clear and close link with health as well as to genetic data. (CoE, 1997: par. 1)

Neither of these definitions are binding (though the first *may* soon be). However it is clear that, by these definitions, biometric data certainly could be considered as relating to health.

The Article 29 Data Protection Working Party suggest that an assessment of the sensitivity of biometric data – and so, by extension, anti-spoofing data – should take into account the context of the processing (A29WP, 2012b: 15). This suggestion, with which we are in full agreement, highlights two important points. Firstly it signals the importance of contextual factors. Secondly, it suggests that categories such as "personal" or "sensitive" data are not fixed hard and fast, but are fluid and flexible.

The upshot for our guidelines is, as above, that because there is scope for differences of opinion as to the status of certain kinds of data, it is necessary to pay significant attention to the contextual factors at play in any given case. Therefore, the guidelines ought to include support in identifying and responding appropriately to relevant contextual factors.

8.2.2. Proportionality[48]

Proportionality is one of the most important data protection principles, yet precisely what it means and how it can be ensured is unclear (Kindt, 2007: 169-170). Proportionality is captured in 95/46/EC by Article 6 (1b), which states that data must be:

> adequate, relevant and not excessive in relation to the purposes for which they are collected and/or further processed.

Thus the principle of proportionality demands an acceptable balance between (i) the level of data processed, and (ii) the importance/value of the purpose of the processing. The use of "level of data" is intended to be neutral as to the sense in which the data is to be measured: in some cases it may be the *volume* of data that is an issue; in other cases it may be the *nature* (e.g. sensitivity) of

[48] Some of the ideas in this section are developed in Rebera et al (2013).

the data that is important (even if there is a relatively small volume processed).

The inclusion of proportionality as a core principle of data protection reflects the recognition that the processing of personal data is inherently risky. That is to say, other things being equal, faced with two possible ways of achieving a certain goal, one of which involves the processing of personal data, and one of which does not (or which involves processing less data), the one which does not involve the processing is preferable. It should, thus, always be demonstrably the case that the personal data processing is a necessary means to a legitimate end. If the processing is not necessary, it means that the end could be achieved by less intrusive or risky means. So proportionality outlaws reckless or reasonably avoidable processing of personal data.[49]

A number of factors must be considered in the balance of proportionality. Of these, several are contextually variant. Rebera et al (2012: 8):

> In the abstract, only cases at the extremes can be conclusively settled: in the majority of real-life cases, no judgement of proportionality can be given without paying close attention to contextual factors.

(For example, we can judge that, say, the processing of health data in order to manage customers' physical access to an ice cream shop is completely disproportionate and unacceptable; we can say that, e.g., responsibly processing fingerprint data to manage access to critical national security infrastructure is most likely proportionate and acceptable.)

The following kinds of contextual factors are liable to be relevant.

- The purpose of the processing; whether the identity management system is for identification or verification; whether the processing is necessary.
- The "general" context of the processing (e.g. is it a national security context, an educational context, a healthcare context, a workplace context, etc.?).
- The people from whom the data will be gathered (e.g. are there any vulnerable groups, e.g. minors?).
- The kind of data to be processed.
- How exactly the data will be gathered; the availability of alternatives to the processing (e.g. are there less intrusive options?); where the data will be gathered.
- How long the data will be retained; specific details regarding the data handling/storage protocols, data security, etc.; whether the data subject holds their data (e.g. on a smartcard they keep with them) or it is centrally stored (e.g. in a database).

[49] One very significant difficulty derives from the lack of consistency in the application of the principle of proportionality across different data protection authorities in Europe. On this see, e.g., A29WP, 2011; A29WP, 2012a; A29WP, 2012b; A29WP, 2012c.

- The level of consent sought from data subjects; the level of information given to data subjects.
- The track record of the data processor (e.g. have they any record of failing to respect data protection frameworks?).
- The level of oversight by (or prior and on-going consultation with) local or national DPAs (data protection authorities).

Proportionality is subject to contextual variation. But contexts are not static but dynamic: they change over time. Accordingly it is necessary to conduct regular reviews to ensure that a system continues to be a necessary means to a legitimate end. As regards spoofing, an important point to consider is this: if an already deployed system is supplemented with anti-spoofing provision, proportionality needs to be re-assessed. It should be re-assessed relative to the overall goal of the biometric system: the proportionality of the anti-spoofing provision should not be judged on its own, in isolation from the wider system, against the security goal of preventing spoofing. Enhancing the security of the system is not in itself an isolated goal, but is a sub-element of the wider goal of ensuring an efficient and effective biometric system.

8.2.3. Consent

The notion of consent applies across all practices entailing ethical risks. Here, we restrict our focus largely to data protection and privacy issues.

95/46/EC defines *consent* as follows:

> "the data subject's consent" shall mean any freely given specific and informed indication of his wishes by which the data subject signifies his agreement to personal data relating to him being processed. (Art. 2(h))

However a European Commission Communication of 2010 reports a number of concerns regarding the concept of *consent* (EC, 2010).

1. The description of *consent* as provided in 95/46/EC is not entirely clear and is sometimes interpreted differently in Member States.
2. Giving informed consent is increasingly difficult online, given the opacity of online privacy policies.
3. In some cases it is unclear what would constitute freely given, specific and informed consent to data processing (e.g. in the case of behavioural advertising, internet browser settings are sometimes but not always considered to imply consent).

Since anti-spoofing provision may involve the processing of personal data, and since it may sometimes involve the processing of sensitive personal data – a form of processing that may

require "explicit consent" (95/46/EC, Art 8(2a)) – it is necessary here to briefly examine some of the main problematic issues arising in this area.

The second and third of the Commission's concerns pertain mainly to the online environment; to the extent that they do we let them pass here.[50] We begin by examining the first, namely that the 95/46/EC definition is unclear.

The 95/46/EC definition uses at least three phrases that could be interpreted in more than one way:

- Freely given consent;
- Specific and informed consent;
- Indication.

8.2.3.1. Freely given consent

Consent is traditionally linked with ideas of *control*, *informational self-determination*, and *autonomy*. Consent must therefore be freely given ("forced consent" is an oxymoron). But a question then arises as to, firstly, what constitutes *freedom* in this context and, secondly, how can that sort of freedom (whatever it is) be reliably recognised? We will discuss the latter question below (§3.2.3.3), in conjunction with the question of how consent in general can be indicated and recognised. In this subsection we look at the question of what it means to *freely give* consent.

In this context *freedom* surely implies both *freedom to make decisions for oneself* and *freedom from coercion*. Let us explore a little more closely what these two constraints entail.

Freedom to make decisions for oneself

The former requires, at a minimum, two things. Firstly, it requires that there should genuinely be an option to give consent and follow one course of action, or to not give consent and follow a different course of action. Both options have to be possible.[51] Moreover, both options have to be *realistically* possible: there should not be any disproportionate disadvantage or unwelcome consequence attending one of the options.[52] Secondly, it requires that one should be sufficiently

[50] This is not to say that biometrics cannot be used online; only that the focus in Tabula Rasa has not been on such cases.

[51] Cf. Locke's (1975[1689]) use of the "locked room" example to illustrate the difference between an action *freely* performed and an action *voluntarily* performed. The man who willingly remains in a locked room (to which he does not have the key) does so voluntarily, but he is not free to leave.

[52] Cf. Hume's (1960[1748]: 156) comment in "Of the Original Contract": "Can we seriously say, that a poor peasant or artisan has a free choice to leave his country, when he knows no foreign language or manners, and lives, from day to day, by the small wages which he acquires? We may as well assert that a man, by remaining in a vessel, freely consents to the dominion of the master; though he was carried on board while asleep, and must leap into the ocean and perish, the moment he leaves her."

well-informed of one's situation as to recognise that a decision is there to be taken. That is, not only should both options be reasonably possible, the subject should recognise that there are two options and that they can decide between them.

This determines the following conditions on freely given consent:

C1. There must genuinely be more than one option, to consent or not to consent.
C2. Both options, to consent or not to consent, must be realistically possible for the subject, and neither should be disproportionately (dis)advantageous.
C3. The subject must be genuinely aware that there is a decision to be taken (to consent or not to consent) and that the decision is their own.

On the face of it at least, C1 rules out the forcible or compulsory gathering of personal data. C2 rules out the use of coercion to force the surrender of personal data (coercion is discussed below). While C3 rules out, on the face of it at least, the covert gathering of data: data subjects ought to be aware that their data is being gathered, and to have the possibility to not consent to giving their data. Of course there may be situations in which these conditions are overridden, e.g. in certain law-enforcement scenarios; but these are to be discussed below.

Freedom from coercion

To *coerce* is to force someone to do something regardless of whether or not they wish to. This means that coercion includes not only those cases in which a subject is made to do something they would rather not do, but also those cases where a subject is compelled to do something which they would, even the absence of such compulsion, have done (or have wished to do). It follows then that coercion is potentially very hard to detect in some cases (namely the latter).

Coercion may take many forms. Covert, forcible, or unavoidable gathering of data involves coercion since the data subject is, in such cases, (respectively) *unaware* of the data gathering, *unable to object to* the data gathering, or *unable to resist* the data gathering. But coercion can be more subtle. For example, where there is a significant power imbalance between the person or agency gathering the data and the data subject – irrespective of whether the data subject consciously feels "under the influence of" or "coerced by" the data gatherer – the data subject's capacity to consent is reduced. This is at least part of the reason why children are not able to give consent, and is certainly a significant factor in the gathering of consent in the arenas of, for example, employment, health, and law enforcement (the adult-child, the employer-employee relation, the doctor-patient relation, and the law enforcement officer-citizen relation all involve significant power imbalances). A great difficulty in this regard is that very many political and social philosophies and analyses see power imbalances as inherent in the structure of society.

A further difficulty facing the view that consent requires an absence of coercion is that it is near

impossible to conceive of a situation in which an individual is not subject to at least some degree of outside influence. We are, after all, social beings who regularly and unavoidably adjust our behaviour and decisions in response to the actions of others. Moreover, if we take a very hard line on the matter, even our own desires or mental states can come to seem as "outside" influences if, for example, they derive to some significant degree from an external source.[53] This is part of reason why consent cannot usually be given under the heavy influence of alcohol or drugs.[54]

It seems plausible to suggest that *complete* freedom from coercion or outside influence is impossible; and while we cannot provide an argument for that claim here, we will adopt it as an assumption. We also adopt, as an un-argued but plausible assumption, the view that consent is often sought and given quite unobjectionably and legitimately. From these two assumptions we draw the conclusion that consent does not require the *absence* of coercion or outside influence. But since these obviously *can* undermine consent, it follows that coercion and influence is not an *all or nothing* affair, but a continuum. At one end of the continuum are levels of influence that are compatible with the giving of consent. At the other end of the continuum are levels of coercion and undue influence incompatible with the giving of consent. In the middle there is a murky area at which closer investigation is necessary to determine whether those levels of coercion and influence are compatible with the giving of consent.

The requirement of freedom from coercion and undue influence therefore entails the following condition on freely given consent:

C4. A subject's stated decision to consent (or not to) must be their own decision. The decision need not be entirely independent of outside influence, but should not be subject to undue influence or coercion.

(Granted the word "undue" carries a heavy burden in C4. This is not however the place to resolve this problem.)

[53] Cf. Kant (1997[1785]) on the idea of a *heteronomous* will.

[54] The relationship between *coercion*, *influence*, and *undue influence* is subtle. It is sometimes argued that coercion differs from forms of influence in that it involves the threat of physical harm to a person or their property (or more generally to something they care about, e.g. their family). It has recently been argued, for example, that undue influence falls short of coercion (threat of harm) because it involves only a "cognitive distortion" to the subject's capacity to balance benefits and risks (Largent, Grady, Miller, Wertheimer, 2012). However the distinction between the physical harm of coercion and the "cognitive distortion" or "mental" or "moral domination" of influence seems somewhat tenuous, at least as far as the capacity to consent is concerned (it may be of far greater significance in legal matters). A decision taken under an outside influence may very well have consequences including physical harm. Moreover a sharp distinction between the physical and mental is not obviously robust (nor either is, for example, the distinction between the physical and cyber or virtual; on this see Joyner & Lotrionte (2001) on information warfare as a form of coercion).

8.2.3.2. Specific and informed consent

The requirement of specific and informed consent is the requirement that the data subject should have a satisfactory understanding of the details of the proposed data gathering/processing, including their associated rights (e.g. rectification of inaccurate data, right to withdraw their data or their consent, etc.). In outline this is clear enough: in detail, i.e. when we examine the specifics of a given case, that clarity begins to blur. Pertinent questions include: how much information do data subjects require? What level of detail is required? Do different data subjects require different levels of information? How can it be ensured that subjects have an adequate understanding of the information they are given? Whose responsibility is it to ensure the subjects' understandings (the subjects or the data processors)?

In some cases it is clear, or at least generally accepted, that a subject is unable to give genuinely informed consent (e.g. minors, unconscious medical patients, etc.). But even here there are difficult cases (there are differences across EU Member States, for example, regarding rules on judging the capacity to consent). Even for those individuals who are routinely considered capable of giving consent, issues may arise. This is particularly so in relation to the information society and rapid technological advances. Not all people understand, for example, what happens to the data that they provide online, or when they open a bank account, or book an international flight. And this is not a matter of the divide between "digital natives" and "digital immigrants" (i.e. those who have grown up amidst digital and online technologies, and those who did not), but a more generalised matter. It is a difficult issue, and one which the developers and end-users of anti-spoofing technologies need to consider.

Hence the requirement of specific and informed consent entails the following condition:

C5. In order to be in a position to give consent, a data subject should understand (or have been given sufficient opportunity to come to understand) the nature and consequences of the action or data processing to which they are potentially subject.

8.2.3.3. The indication and recognition of consent

What constitutes an indication of consent will vary across contexts. In some cases – medical trials for example – witnessed written consent may be required. In other cases, a simple action like nodding one's head – or even a non-action like simply not removing oneself from a (physical or online) space – may suffice. Subjects should therefore be made aware of how the data controller registers consent: genuinely informed consent cannot be tacit.

The appropriate manner of indicating consent may also vary with the kind of data that is at stake. In relation to personal data, *consent* is taken as a "freely given specific and informed indication" of the data subject's wishes (95/46/EC, Art. 2(h)). However, the processing of sensitive data may

require "explicit consent" (95/46/EC, Art 8(2a)). Explicit consent is normally given in writing because it may then also serve as evidence of consent (A29WP, 2011: 25). It can also be given orally, however, following the Article 29 line, it ought to be somehow recorded (e.g. on video).

With respect to the indication and recognition of consent, we can add the following requirements:

C6. The data processor's means of recognising consent should be clear to the data subject. The danger of the subject "inadvertently" giving consent should be minimised.
C7. There should be a record of the giving of explicit consent (e.g. in writing).

8.2.3.4. Summary

The table below indicates how the seven implications of the 95/46/EC definition of consent might be relevant to the deployment of anti-spoofing technologies. It should be remembered that the requirement to get the consent of data subjects will apply not with specific reference to the anti-spoofing aspects of the technology, but to the biometric system as a whole. Nonetheless, the table considers only anti-spoofing provision.

C1.	**There must genuinely be more than one option, to consent or not to consent.**
	It does not follow from this principle that, when a data subject has no opportunity to object or to not consent to the data processing, the data processing is therefore illegitimate. But it *does* follow that consent cannot, in such circumstances, be legitimately given as the legal ground of the processing. For example, the blanket monitoring, by way of an intrinsically robust biometric such as gait analysis, of a public space which can hardly be avoided in the course of normal day-to-day life, cannot be justified by appeal to consent: in such a case the option to not-consent does not exist (and hence, logically, neither does the option to consent).
C2.	**Both options, to consent or not to consent, must be realistically possible for the subject, and neither should be disproportionately (dis)advantageous.**
	A subject's capacity to provide meaningful consent is compromised when the consequences of consenting are very significantly better or worse than the consequences of not consenting. So if the subject is faced with the choice of either (a) going through international border control in 2 minutes by providing their fingerprint, subject to liveness detection, or of (b) going through international border control in 3 hours by queuing for the manual paper passport control, then any consent they provide for the fingerprint/liveness detection is highly questionable. Consent would, in such circumstances, be an unreliable ground of the data processing.
C3.	**The subject must be genuinely aware that there is a decision to be taken (to consent or not to consent) and that the decision is their own.**
	Imagine a situation in which some form of physical access (say, to a workplace) is managed by a facial recognition system with no specific anti-spoofing provision. Imagine further that the employees have provided satisfactory consent to the processing of their data in this way. If the employer replaces their facial recognition system with a new system which incorporates anti-spoofing technology which implies further or different data processing, the employer is bound

	to inform the employees of the change and of the fact that they are able to refuse consent to the data processing if they so wish. If the employer does not, the employees may be unaware of the change and unaware that there is a new decision to be taken.
C4.	**A subject's stated decision to consent (or not to) must be their own decision. The decision need not be entirely independent of outside influence, but should not be subject to undue influence or coercion.**
	In the deployment of anti-spoofing technologies (or biometric systems utilising them), there may be any number of reasons why a data subject is subject to undue influence or coercion. Insofar as these stem from power imbalances or other reasons deriving from the *specific context* of deployment, we cannot fruitfully discuss them here. We can, however say the following. A data subject faced with the choice of consenting or not consenting to their personal data being processed by an anti-spoofing technology should *not* be basing their decision on the assumption that such technologies are *only* introduced because serious risks and threats are entertained in their absence. Or at least if they *do* base their decision on a skewed view of the risks and benefits of giving or not giving their consent, that skewed view should not have been brought about by a rhetoric or narrative of fear, suspicion, risk, and insecurity surrounding the technology itself. To put it another way: the *mere presence* of an anti-spoofing technology should not itself influence a data subject's decision to consent to their data being processed by means of it. Such a situation would occur if the subject's thought process was something like this: *since they have anti-spoofing provision it must be necessary (else why would they have it?), therefore I had better consent*.
C5.	**In order to be in a position to give consent, a data subject should understand (or have been given sufficient opportunity to come to understand) the nature and consequences of the action or data processing to which they are potentially subject.**
	Broadly, this requirement gives rise to two main difficulties: (i) explaining complex or technical information in an easily comprehensible manner; and (ii) communicating a perhaps quite large amount of information in such a way that the process is neither unduly disruptive[55] nor boring.[56] Regarding anti-spoofing provision in particular (as opposed to biometrics in particular), the main problem here concerns (ii). Biometric systems with anti-spoofing provision will naturally tend to require the successful communication of a larger amount of information. The challenge is to convey that information – e.g. what anti-spoofing measures are in place, what data they process, and why they are necessary – without exerting undue influence on the data subject (cf. C4 above), and without compromising system security (for example by giving away so much, or so crucial, information as to aid spoofers.
C6.	**The data processor's means of recognising consent should be clear to the data subject. The danger of the subject "inadvertently" giving consent should be minimised.**
	One respect in which this requirement might have specific importance is in those cases where

[55] To take an extreme example, for an ATM using multimodality and a variety of liveness detection techniques to explain exactly what data it takes, how it uses it, whether and if the data is retained, transferred (and so on…) could take a long time, causing queues and delays sufficient to counteract the convenience that the ATM was introduced to provide in the first place. This would be disruptive.

[56] This is one of the problems with online privacy policies and (and online and offline) terms and conditions: they are often an unreliable method of ensuring *informed* consent because most people (quite reasonably) cannot be bothered to read them.

	an existing biometric system is upgraded to add anti-spoofing provision. Once a data subject has given consent to the data processing implied by a given biometric system, their continuing use of the system is reasonably taken as an indication that they have not revoked their consent. When the system is changed or upgraded, it is no longer acceptable for the data processor to take the data subject's continued use of the system as an indication of consent. A *new* indication of consent is required and the data subject must be informed that this is the case.
C7.	**There should be a record of the giving of explicit consent (e.g. in writing).**
	This requirement has no specific implications for anti-spoofing as opposed to biometrics in general. However it is worth noting that the question of whether certain biometric data constitutes health – and so *sensitive* – data, if it is pertinent at all, is particularly likely to be pertinent in the case of inherently robust biometrics such as gait, EEG, or ECG. In such cases, if the data is judged sensitive, explicit consent will be required.

Table 9. Impact of consent on anti-spoofing provision

8.3. European Data Protection in the Near-future

The aim of this final subsection of Part 8 is to very briefly discuss some of the proposed changes to the European data protection framework. This discussion is necessarily speculative, and consequently brief. But to the extent that the changes currently under discussion reflect perceived gaps in the existing European data protection framework – and to the extent that these gaps are gaps in the protection and promotion of fundamental rights – we may hope to discern from them pointers towards best-practice.

We wish to highlight three points from the Commission's Proposal for a General Data Protection Regulation (EC, 2012), concerning: the definition of "biometric data"; the indication and proof of consent; and principles concerning the informing of data subjects.

8.3.1. "Biometric Data"

The proposed Regulation differs from 95/46/EC by offering an official definition of "biometric data".

> "biometric data" means any data relating to the physical, physiological or behavioural characteristics of an individual which allow their unique identification, such as facial images, or dactyloscopic data. (EC, 2012: 42)

It should be noted that the Article 29 Data Protection Working Party (2102c: 10) have recommended that the definition be re-worded as follows:

> *"biometric data" means any data relating to the physical, physiological or behavioural*

> *characteristics of an individual which* **are unique for each individual specifically**, *such as facial images, or dactyloscopic data.*

The pertinent issue for present purposes is to establish any consequences of this definition for the question of whether biometric data constitutes personal and/or sensitive data.

Under the new proposals, "personal data" will be defined as any information relating to a data subject (EC, 2012: 41). A "data subject" is defined as:

> an identified natural person or a natural person who can be identified, directly or indirectly, by means reasonably likely to be used by the controller or by any other natural or legal person, in particular by reference to an identification number, location data, online identifier or to one or more factors specific to the **physical, physiological, genetic, mental**, economic, cultural or social identity of that person. (EC, 2012: 41; boldface added)

In conjunction with the definition of "biometric data", the factors highlighted in boldface make clear that, as it currently stands, biometric data is likely to continue to be considered personal data.

With regard to the question of whether biometric data is *sensitive* data, we can refer now to the proposed definition of "data concerning health":

> "data concerning health" means any information which relates to the physical or mental health of an individual, or to the provision of health services to the individual. (EC, 2012; 42)

However the difficulty here is that the definition itself uses the concept of "health". So this question, at least as far as the new data protection proposals are concerned, must remain open.

8.3.2. On Consent

The proposed Regulation makes clear that consent must be *affirmative*, i.e. actively signalled.

> Consent should be given explicitly by any appropriate method enabling a freely given specific and informed indication of the data subject's wishes, either **by a statement or by a clear affirmative action by the data subject**, ensuring that individuals are aware that they give their consent to the processing of personal data, including by ticking a box when visiting an Internet website or by any other statement or conduct which clearly indicates in this context the data subject's acceptance of the proposed processing of their personal data. **Silence or inactivity should therefore not constitute consent**. Consent should cover all

> processing activities carried out for the same purpose or purposes. If the data subject's consent is to be given following an electronic request, the request must be clear, concise and not unnecessarily disruptive to the use of the service for which it is provided. (EC, 2012: 21; boldface added).

Secondly, the proposed Regulation makes clear that the burden of proof for the data subject's consent lies with the data controller (EC, 2012: 45).[57]

These two points reinforce the suggestion made above that consent should be unambiguous and clearly recorded.

8.3.3. Informing Data Subjects

Three aspects of the Preamble to the proposed Regulation relate to the obligation to inform data subjects of the processing (EC, 2012: 25).

Preamble (48) states that the data subject should be informed of the "the existence of the processing operation and its purposes, how long the data will be stored, on the existence of the right of access, rectification or erasure and on the right to lodge a complaint" as a matter of "fair and transparent processing" (EC, 2012: 25). Preamble (49) determines that this information should be provided to the data subject at the time of collection (or, if the data are not collected directly from the subject, within a "reasonable" period). Data subjects should also be informed on the first occasion when their data is legitimately disclosed to another recipient. Preamble (50) sets some limitations on the previous requirements, including that these obligations on the controller are not to be imposed "where the data subject already disposes of this information, or where the recording or disclosure of the data is expressly laid down by law, or where the provision of information to the data subject proves impossible or would involve disproportionate efforts" (EC, 2012: 25).

These obligations are quite comprehensive, meaning there is a lot of information to be conveyed to the data subject. There could, in principle, be some appeal to the final caveat, i.e. that the information need not be provided where this "proves impossible or would involve disproportionate efforts". But it should be considered firstly that although this clause offers some possibility of escaping the obligation, it could certainly be interpreted strictly so that the vast majority of deployments of biometric systems *did* have to provide the information; and secondly, it should be considered that, independently of how this clause is interpreted, and even independently of whether the clause is included in the final Regulation, the rationale behind it is

[57] A third important point, which we will not stop to consider at greater length than this footnote, is that the Regulation proposes that "Consent shall not provide a legal basis for the processing, where there is a significant imbalance between the position of the data subject and the controller" (EC, 2012: 45). This is in accordance with the discussion above in Subsection §3.2.3.1.

quite clear: other things being equal, data subjects have a right to know *that*, *why*, *how*, and *by whom* their data is being processed. The guidelines we provide in this document therefore respect this rationale as matter of good ethical practice, quite independently of its provision in any current or future data protection legislation.

8.4. Impact Assessments

Impact assessments prior to technology deployments are now relatively common. In this document we consider: *privacy* (PIA), *data protection* (DPIA), and *ethical* impact assessment (EIA).

Details of these impact assessments differ from case to case. For instance, some PIAs will incorporate DPIAs. Yet since the European fundamental rights framework distinguishes privacy from data protection, it can make sense to analyse impact on these rights separately (De Hert, 2012). Wright and De Hert (2012: 8-9) employ the following definitions.

> PIA: "a process for identifying and evaluating risks to privacy, checking for compliance with privacy legislation and considering ways in which those risks can be avoided or mitigated".
>
> DPIA: "primarily a compliance check [with relevant data protection legislation]".

EIA is a newer concept than PIA, but is gaining a foothold in the debate on the management and introduction of (especially new and emerging) technologies. EIA is, broadly:

> A framework that could be used by those developing new technologies, services, projects, policies or programmes as a way to ensure that their ethical implications are adequately examined by stakeholders before possible deployment and so that mitigating measures can be taken as necessary. (Wright & Mordini, 2012: 416)

There are many differences between impact assessments as carried out by different organisations in different domains in different parts of the world. Although impact assessments are not always mandatory, it seems clear that they can offer significant benefits to all stakeholders. In the Tabula Rasa ethical guidelines, we have recommended that impact assessments for privacy, data protection, and fundamental rights be carried out prior to the deployment of anti-spoofing technologies. We have not provided templates for these, but have given some overall guidelines on identifying key factors (§5.4, §5.5), including the identification of contextually salient factors (§5.6).

9. Annex 2 – Ethical Issues

Research in WP7 has turned up several ethical issues, which are presented in this section. We organise the Subsection around four thematic areas:

- Data protection, privacy, and fundamental rights;
- Spoofing, deception, and honesty;
- Trust, transparency, and secrecy;
- Societal factors: living with biometrics and anti-spoofing.

9.1. Data protection, privacy, and fundamental rights

Since data protection is discussed above (§§3, 5, 8), we say only a little about it here.

It goes without saying that high standards of data protection are essential. One reason is that personal data is increasingly economically valuable—an issue with which society has arguably not really yet come to terms. The role of the individual data subject in the personal data marketplace remains unclear. Do the benefits one gains from participating in a supermarket's loyalty card scheme genuinely balance the value of the consumer habits information that the supermarket gains? Is the benefit of connecting with friends on Facebook really worth the personal data that is given up in return for the service (or might a less intrusive social network be better "value"?)? It is not clear, and, as yet, there is no clear and consistently applied means of assessing the value of personal data.

This is a general point about personal data, but it brings into focus the need to ensure that anti-spoofing technologies utilise and store the smallest amount of data consistent with functionality. Data minimisation is also important because it reduces the risk of function creep (i.e. the use of a system or the data it processes for purposes beyond those originally intended; on this see Mordini & Massari (2008)).

To avoid function creep, it can be necessary to build "tight" systems, i.e. systems which avoid redundancy wherever possible. A tight system gathers and retains no more data than is absolutely necessary, and includes no functionality that is either under-utilised or not utilised at all. However, one of the problems of spoofing identified in the course of WP7 research is that we have imperfect knowledge of the current vulnerabilities of biometric systems. Moreover, we have extremely imperfect knowledge of the *future* vulnerabilities of biometric systems. This problem is an aspect of what, in D7.2, we christened "the quantification problem". Now a natural response to this problem is to attempt to develop systems which intentionally include a certain degree of "slack", i.e. which have a certain flexibility which means that they can be tweaked and adapted in the

future. In this way, vulnerabilities that emerge can be responded to by developers, obviating the economically and organisationally unwelcome need to replace already deployed systems. But there is a conflict here. On the one hand we have an imperative to create rigid, tight systems in order to minimise the risk of function creep. On the other hand, we have an imperative to create flexible systems to promote responsiveness to emerging vulnerabilities. On both sides we have an admirable and desirable goal, and yet the two are, on the face of it, incompatible.

A plausible answer to this problem is to create flexible systems as per the second option, but to do so amid a culture in which the risks of function creep are minimised in other ways. These may include promoting greater stakeholder awareness of function creep risks, promoting awareness and monitoring (e.g. by Member State national data protection authorities), and promoting a generalised practice of openness and transparency in biometrics and anti-spoofing. The first two of these methods are general responsibilities; the third is connected with the goals of openness, transparency, and "active honesty" in relation to anti-spoofing technology (discussed below, §9.2).

Creating flexible anti-spoofing systems need not involve gathering redundant data—indeed it *should not*. So data minimisation is a goal no matter what. The gathering of data is of concern also with respect to other fundamental rights (cf. Table 9, §8.1).

Existing critiques often focus on the threat to dignity and integrity from the way in which biometrics conceptualises the human body.[58] Such critiques often presuppose – understandably – that there is a commonly accepted, robust understanding of what dignity is. This assumption, as was pointed out by Anders Sandberg in his Tabula Rasa interview (D7.2), is questionable. Notwithstanding the fact that dignity is one of (if not *the*) core values of the European ethical framework, the question of what exactly it *is* has no simple answer. The same may be said of *integrity* and – not least – *privacy*, regarding which there is an enormous literature.[59] [60]

Although a precise account of *dignity* is lacking, we are not completely in the dark as to its nature. In the European framework – indeed in the United Nations framework – dignity is considered to

[58] E.g. Agamben (2008) or various works of Irma Van der Ploeg (1999; 2002; 2007), for instance. For a different view of the conceptualisation of the biometric body see Mordini & Rebera (2012).

[59] Although it does not technically begin here, a good starting point for delving into the privacy literature is Westin (1967). Influential more recent book-length treatments include Regan (1995), Etzioni (1999), Solove (2008), Nissenbaum (2010), *and many more*.

[60] Part of the difficulty is that concepts such as *dignity*, *integrity*, *privacy*, *autonomy* (and so on) are closely interwoven and, most probably, overlapping. Take *dignity*, *integrity*, and *privacy* for instance: scholars tend to distinguish two approaches to the concept of integrity. According to Mordini, the first approach "contends that the right to be free from bodily (and mental) intrusion is inherently part of the notion of human dignity" (Mordini, 2011: 192); while the second approach "maintains that bodily integrity is the right of "every human being ... to determine what shall be done with his own body" [fn: Schloendorff v. Society of New York Hospital, 1914, quoted by Maschke (2003)] and to protect his physical privacy" (Mordini, 2011: 192). If this is right, the concept of *integrity* is bound up (*tangled*, perhaps) with the concepts of *dignity* and *privacy*. This suggests that one cannot fully grasp any of these key ethical concepts without a firm grip on all of the others. So the difficulty is severe and on-going.

be a foundational ethical concept. The "Explanations Relating to the Charter of Fundamental Rights" (EU, 2007) make this explicit:

> The dignity of the human person is not only a fundamental right in itself but constitutes **the real basis of fundamental rights**. The 1948 Universal Declaration of Human Rights enshrined human dignity in its preamble: "Whereas recognition of the inherent dignity and of the equal and inalienable rights of all members of the human family is **the foundation of freedom, justice and peace in the world**." (EU, 2007: 17; boldface added)

Here we have one reason why affronts to integrity, autonomy, privacy (and so on) are often interpretable as affronts to dignity.[61] Avoiding, then, the task of adequately distinguishing dignity, integrity, autonomy, and privacy, we wish to focus on two ethical issues raised by anti-spoofing in connection with dignity broadly construed.

9.1.1. Gathering of "too much" or "intimate" data

Gathering personal data can be problematic in at least two ways. Firstly, it can be problematic if too much information is gathered. Individuals do not, in general, enjoy the feeling that their privacy has been breached, or that someone or some organisation knows a very great deal about them, or that they are being monitored. Moreover, they have, in many situations, an ethical and legal right not to be in these situations. Anti-spoofing technology must not, however, gather data that is not absolutely necessary for its functioning.

In Tabula Rasa, we have addressed this point through the recommendation that redundant data not be gathered, and that it be destroyed if it is gathered (D7.3). Furthermore, the coordinator, Sebastien Marcel, confirmed in D7.2 that the anti-spoofing techniques developed in Tabula Rasa rely on data gathered anyway in the normal functioning of the biometric system in question.

Secondly, the gathering of personal data can be problematic if the data is sensitive or intimate. In D7.2 we raised the issue of "intimate data", i.e. data which, in some sense, gets to the core of who and how one fundamentally is. Examples include information about personality, mood, or current emotional state (D7.2, pp. 56-7; cf. also Rebera et al, 2013). This kind of information is not at stake in Tabula Rasa. Nevertheless, it has anti-spoofing applications (it could be used to detect

[61] To be precise: what we have here is a reason why affronts to integrity, autonomy, privacy (and so on) may *appear to be* affronts to dignity. It could be the case that dignity alone is the fundamental right underpinning all others. It could alternatively be the case that dignity is a basic fundamental right alongside some others (integrity and autonomy perhaps). We haven't considered any arguments for either position here. On either account, there are ready explanations of why affronts to other values can often appear as affronts to dignity. If dignity alone is fundamental, then insofar as it underpins other rights, affronts to those rights are indirect affronts to it. Alternatively, if dignity is only one of many fundamental values, affronts to the others may be easily confounded with affronts to dignity since both have the same effect: they undermine one's moral status as a bearer of certain inalienable rights.

coercion attacks, e.g.). The intimacy of this kind of data needs to be taken very seriously. The fact that it may not, from a technical or legal point of view, be strictly "sensitive" should not lead to its being side-lined in assessments of proportionality.

The connection of personal data – sensitive, intimate, or otherwise – with our sense of who we are is grounded in the fact that identity is a very complex thing. A recent report for the UK Government Office for Science undertakes to speak only of "identities" in the plural, rather than "identity" in the singular, and makes the following points.

- Identities refer to the way in which individuals perceive themselves and their place in the world, and how they are categorised by others. People have coexisting, multifaceted, overlapping identities, which alter in emphasis depending upon the context. At home a person might find their identity as a parent most important, while at work they might identify as a company employee, and online pursue a hobby as part of an interest group. Some identities can change over the life course of an individual.
- Identities can be elective or chosen by an individual, such as by self-categorised membership of a social group, or ascribed and controlled by others, for example through data held on a person by a supermarket.
- Identities can be inclusive, such as membership of a family, team, religion or other group, or exclusive, defined by not being a member of a particular group. Again, these identities can be controlled by the individual or by others, for instance through rejection from a group. (All three bullet points from: F'sight, 2013: 9)

Each of these points adverts to the fact that identities can be imposed, influenced, or controlled by others. It is thus important that anti-spoofing technology does not contribute to the imposition of negative self-conceptions on data subjects. This might happen in two ways.

The first relates to the way in which one's self-conception might be influenced by the proliferation of biometric systems in general. In this relation, biometrics is subject to the critique that there is something "dehumanising" about it—that it does violence to human dignity in some manner, or that it objectifies individuals or that it reduces them to "bare life". However this broad position was countered, convincingly, by Luciano Floridi (D7.2).

The second way in which anti-spoofing technology might contribute to the imposition of negative self-conceptions is through the lingering association of biometrics (notably fingerprinting) with criminality, and – relatedly – of anti-spoofing with deviancy and suspicion.

The connection of spoofing with deviancy, suspicion, and suspicious activities and individuals has been one of the core findings of WP7 research. It has become apparent – from a variety of perspectives – that spoofing is not always morally wrong. We will shortly discuss this matter

below. For the moment we will turn to a related issue affecting dignity and integrity: the right to anonymity.

9.1.2. A legitimate right to anonymity

In some cases a person will be subject to identification by a biometric system without their knowledge, or in a context in which their awareness of the identification process is compromised (e.g. because the process is so routine that they tend to ignore it or not notice it). Such cases are comparable with certain uses of CCTV: sometimes one enters an area of CCTV either in ignorance of the surveillance, or with only the vaguest realisation of having so done—and one can easily imagine similar scenarios involving, e.g., inherently robust modalities such as gait, or anti-spoofing-equipped face recognition cameras. Such scenarios can be perfectly legitimate. Nonetheless, people have a legitimate interest in anonymity in various situations; and this interest can be compromised by covert or low-visibility gathering of data.

Anonymity has traditionally been associated with *namelessness*.[62] However, as Nissenbaum (1999) points out, namelessness is not the goal of anonymity, but only the *means of achieving* the goal of anonymity. The goal of anonymity is "the possibility of acting or participating while remaining out of reach", where being *unreachable* means, in effect, being practically unaccountable for one's actions, statements, or opinions (Nissenbaum, 1999).

Anonymity is a controversial concept because it is seen as providing a veil behind which unpleasant or criminal acts can take place (e.g. internet trolling)—a "barrier" to accountability (Bynum, 1997). However, in certain circumstances, anonymity is desirable – and necessary – for entirely legitimate reasons. Gary Marx suggests 15 rationales for anonymity (Marx, 1999), which we may compact thus:

	Rationale	Examples
1	To facilitate and/or encourage communication and informational exchanges on matters of public interest.	- Anonymity for whistle-blowers or journalists; - Witness anonymity programmes; - Anonymity for those seeking "sensitive" information or advice (e.g. on sexual health).
2	To facilitate and/or encourage exchanges of resources which might otherwise incur undesirable consequence (e.g. criminal investigation).	- Amnesties on handing in weapons to police; - Needle-exchange programmes; - Sperm/egg donation; - Political/charitable donations.
3	For economic purposes.	- Paying in cash to avoid consumer profiling. - Avoiding identity theft.

[62] This reflects the word's etymology (it comes from the Greek "*anonymos*" meaning "without a name").

4	Rituals, traditions, and play.	- Masked balls, the Venetian tradition of *Bauta*, etc. - Children's games, online role-playing games, etc.
5	Avoiding persecution/discrimination, or encouraging attention to the *content* of what is said or done, rather than on *who says or does* it.	- Victims of repressive regimes; - Academic peer-review; - Applicants for positions (jobs, university places, etc.) having their names (genders, ethnicities, etc.) hidden.
6	To avoid unwanted intrusion in one's time, space, or person.	- Interest in being "left alone" (e.g. not stared at, not asked intimate or forward questions, etc.). - Traditional or customary expectations of anonymity (e.g. opposition in the USA to "caller-ID" schemes). - Unlisted phone numbers to avoid telemarketers. - Assumed names/email addresses when, e.g., signing up for online services. - Celebrities wearing dark glasses, caps, etc. to avoid being recognised.
7	To encourage experimentation, non-conformity, risk-taking (etc.) without undesirable significant consequences (e.g. embarrassment).	- Trying out new identities or activities seemingly in tension with one's "normal" identity or activities (e.g. sexual behaviour, political views, etc.)

Table 9. Rationales of anonymity (derived from Marx (1999)). [63]

Although no article of any document in the European fundamental rights framework specifically sets out a right to anonymity, it arguably falls within the scope of other provisions, including the rights to privacy and dignity.[64]

As far as anti-spoofing goes, two main lessons should be drawn from this discussion of anonymity. The first concerns system deployment. Biometric systems involving the covert gathering of data (e.g. of gait data) should be deployed only when the security gains are appropriately balanced against the threat to fundamental rights and ethical values—in this case, against the threat to, inter alia, individuals' legitimate interests in anonymity. This is an issue that will likely be decided at a relatively high level, i.e. by security professionals, politicians, or senior members of private organisations; it is to be hoped that such decisions can be taken in

[63] It should be said that Marx also sets out 10 rationales in favour of identifiability (i.e. against anonymity). His general conclusion is that what is needed in questions of anonymity (he is concerned in the 1999 paper with anonymity in communication) is a close examination of when anonymity is appropriate and when not, and also an appeal to fairness and honesty. In practical terms, this means that when anonymity is preserved through the use of pseudonyms, the fact that a pseudonym is being used ought to be made plain.

[64] When privacy is construed in terms of "restricted access", anonymity is explicitly taken to be a form of privacy. And some of the examples given in line 6 of Table 9 above may be readily construed as issues of dignity or integrity.

consultation with relevant authorities and/or representatives of salient interest groups (e.g. privacy advocacy groups).

The second lesson concerns anonymisation. Wherever possible, anti-spoofing technologies should include, as standard, all reasonable measures to ensure the data subject's anonymity. Anonymity is not an all-or-nothing concept: there is no such thing as complete anonymity. Nobody seeks complete anonymity, but only *relative* anonymity. In some situations one desires more anonymity: in some situations less. Anti-spoofing technologies can be more acceptable (or perhaps less unacceptable) by not aggravating the problem by removing people's control over their anonymity. In many situations, if people do not want to be identified, they should not be identified without good reason. But even when there is good reason, measures can be implemented to preserve relative anonymity. For example, if covert identification manages or monitors access to a secure facility, it may be necessary to identify an individual *only* as they enter the facility: it does not automatically follow that their identity needs to be stored or retained by the system thereafter.[65]

9.2. Spoofing, deception, and honesty

The conceptual link between spoofing and deviancy is stronger than that between biometrics and deviancy. The biometrics-deviancy connection is historically contingent, having arisen from – for the most part – law and order protocols (fingerprinting, mug-shots, etc.). The spoofing- deviancy connection springs from the fact that spoofing is inherently deceptive. However, it does not follow that spoofing is always morally impermissible (see D7.2). Rebera et al (2013) put it like this:

> [When] spoofing occurs against the backdrop of a situation in which the spoofer's legitimate interests have not been adequately respected, [it would appear that] the spoofer has some license to take steps that disrespect or undermine the authority's autonomy.

The relation between spoofing, deception, and honesty manifests itself in several ways. One is in the slide from the (correct) view of *spoofing as deceptive* to the (incorrect) view of *spoofing as inherently deviant and immoral*. An offshoot of this observation is that it becomes necessary to consider the question of *who spoofs and why*. Or to put it more pertinently, we have to consider the question: *under what circumstances could it be morally justifiable to spoof a biometric system*?

There are three broad contexts in which spoofing a biometric system could be justifiable (Rebera et al, 2013):

[65] Similarly, it could be that a person identified as *not being on a blacklist*, no longer needs their identity stored.

1. The system may, even if it is deployed by a legitimate authority for a legitimate purpose, gather so much (or so intimate) data as to be intrusive (i.e. a threat to one's privacy).
2. The system may be deployed by a legitimate authority for an illegitimate purpose.
3. The system may be deployed by an illegitimate authority.

We are discussing the European framework in this document, hence we discount the third context.[66] We discount also the first as it is discussed above (§9.1.1). Thus our original question boils down to this: *under what circumstances could a legitimate authority's deployment of a biometric system equipped with anti-spoofing provision be illegitimate, i.e. ethically unacceptable*? Painting in broad strokes, the answer is: a legitimate authority's deployment of an anti-spoofing technology is ethically unacceptable when that system infringes, for no good reason, people's fundamental rights and legitimate interests.

Findings of WP7 include: (i) that spoofing is not always unjustifiable; but (ii), the promotion of ethical best practice in the development and deployment of anti-spoofing techniques should, if successful, have the consequence that the likelihood of unjustifiable and unethical acts of spoofing decreases (since the likelihood of illegitimate deployments of anti-spoofing technologies should decrease).

One way of avoiding illegitimate deployments is to encourage "active honesty". One assumes that actors deploying biometric systems are honest about their reasons for so doing and the details of the deployment: honest, that is, in the sense that they *do not lie* about it. But further, their honesty should be *active*, i.e. that should actively take steps to inform relevant stakeholders. Relevant stakeholders may include the general public, some subset of the public (e.g. a particular company's employees), or a watchdog of some kind (e.g. a local data protection authority).

Honesty – and, indeed, active honesty – should not be confused with *openness*. Active honesty does not entail broadcasting details to everyone. Rather, it entails informing the *appropriate stakeholders*. Sometimes the appropriate stakeholder will be the general public. But often, details of a deployment will be of interest only to a national data protection authority. Similarly, it may in many cases be appropriate to tailor information to specific stakeholders. For example the public may need to be informed that face-recognition-equipped cameras have been installed in their town centre for security purposes; they may need to know some details of how their personal data will be handled. However, the local data protection authority will need to know many more details. Here – as elsewhere in this domain – the details of particular cases are context-dependent.

[66] Which is not to deny that questions have been raised as to the "democratic deficit" in Europe (Bellamy, 2012; Bonde, 2011), particularly in relation to the current economic situation (Motha, 2012; Nicolaïdis, 2013).

9.3. Trust, transparency, and secrecy

Transparency and secrecy are not polar opposites. The formal complement of transparency is opacity, not secrecy; and opacity differs from secrecy in that the latter involves *intention*. To hold something secret is to intentionally conceal it (Bok, 1989). Thus when we speak of transparency and secrecy in relation to anti-spoofing, we have two thoughts in mind. First, in promoting transparency, what is promoted is the clearing away of obstructions barring from view the structures by and through which anti-spoofing technologies are developed and deployed. Second, in advocating against secrecy, what is promoted is a culture, or an ethos, whereby intentional concealment is the exception rather than the rule (we do not suggest a complete lack of secrecy, for it may be necessary in certain circumstances). This second thought has the intention of bringing into the light as much information as is possible (i.e. as is possible bearing in mind security considerations); the first has the intention of making what is in the light clearer to view.

The major transparency-related tension unearthed in WP7 has been the question of how open system developers ought to be regarding vulnerabilities to spoofing attacks. The following factors are significant.

- Public trust and confidence in biometrics, anti-spoofing, and ICT in general, needs to be encouraged by developing and publicly promoting secure systems.
- The "quantification problem": it is very difficult to accurately measure vulnerabilities. Partly this is due to contextual factors. Hence while it is one thing to be open about vulnerabilities, it may be quite another to be *accurate* about them. Accuracy is not always easily achieved.
- Openness about operational procedures must accompany openness about technologies.
- How much information about system vulnerabilities should be made available, and to whom should it be made available? *Transparency* does not entail *complete publicity*.
- Information disseminated needs to be honest and realistic.
- Information is only helpful if its target audience can understand it. Setting out technical details in comprehensible form is a challenge.
- Vendors of biometric systems may have an economic interest in not revealing vulnerabilities. Approaches to overcome this reluctance may include: legislation; industry standards and codes of conduct; and economic incentives.

Let us turn to the central question in this area: how much information should system vendors give out regarding vulnerabilities, and to whom should they make it available?

A general rule of thumb would appear to be this. ***Vendors should make available to customers sufficient information regarding security and vulnerabilities that the customer is able to make a well-informed judgment as to whether the system in question is good fit for the specific***

context(s) in which they intend to deploy it. But what does this mean in more specific terms?

What kind of information would enable a customer to make a well-informed judgment as to a system's suitability? The customer would need to come to terms with two factors. First, what are the main requirements engendered by the specific deployment contexts they are contemplating? Second, do the technical specifications of a given system meet those requirements?

Responsibility is shared between vendor and customer. In relation to the first factor, the customer must make an accurate analysis of the deployment context they are considering. The second factor concerns the possibility of the customer understanding the information that the vendor provides. Here responsibility lies with the vendor to provide sufficient information in a comprehensible form. "Sufficient information" means information that is accurate and comprehensive enough that the customer can, for each of the main factors deriving from his use context, assess the likely performance of the system in question.

In the abstract we cannot specify precisely which information should be confidential. Thus we offer again only a rule of thumb. ***Other things being equal, vendors should be completely open about the vulnerabilities of the systems they sell to spoofing attacks***, i.e. if people ask, they should tell.

As against this, one could argue that vendors should not aid spoofers by publicising vulnerabilities. However the following response appears compelling. Suppose it to be true that, if spoofers are aware of a technology's vulnerabilities, they gain an advantage. Given that that there are more ways of finding out about vulnerabilities than simply being told by a vendor (e.g. the spoofer could buy the technology himself and experiment with it), the likelihood of completely suppressing information about vulnerabilities is very low. It is thus better to be open about vulnerabilities and to attempt to develop other responses to them (e.g. procedural standards, human oversight, etc.).

Unlike vendors, system developers may justifiably hold information back. A young technology which gains a reputation for insecurity may not even make it to the market, or may find its uptake curtailed. This could be a significant loss to overall security in the long term. Assuming that we are talking about technologies not yet in the marketplace, there is little or no harm in suppressing vulnerability information since nobody is actually going to deploy the technology in ignorance of its vulnerabilities and thus suffer the consequences of its being spoofed.[67] Thus, ***information concerning the vulnerabilities of pre-deployment stage technologies may be held confidential, if releasing it would compromise the likelihood of those technologies reaching deployment***.

[67] One might argue that some harm could accrue through the opportunity cost of not sharing research findings among the whole research community. We leave this issue to one side as it is not the main line of research in Tabula Rasa.

9.4. Societal factors: living with biometrics and anti-spoofing

In D7.2 we discussed the possible impact of the widespread deployment of anti-spoofing technologies on how we live: what it is (or might be) like to live in a world where biometrics and anti-spoofing technologies are more common and more normal than they are at present. The main issues included *privacy and trust*, *the surveillance society*, *personal identity*, people's *conceptions of the self*, and the *public perception and societal impact of secure biometrics* (D7.2, §§3.6-3.9).

Here we wish to emphasise the practical importance of such issues for the guidelines in this document. ***Privacy, and respect for the fundamental rights of data subjects, need to be built-in to anti-spoofing technologies***. As we have indicated, context is extremely important in assessing the ethical implications of a given deployment. But for developers, who cannot work with any one specific deployment context in mind, it is important to have some generic standards. Thus the guidelines are intended to support developers in developing ***anti-spoofing technologies which, when deployed, do not from their general design or functionality cause or aggravate ethical issues***.

[end]