Privacy Through Anonymisation in Large-scale Socio-technical Systems: Multi-lingual Contact Centres across the EU

Privacy through Anonymisation in Large-scale Socio-technical Systems: Multi-lingual Contact Centres across the EU

Claudia Cevenini, Enrico Denti, Andrea Omicini, and Italo Cerno

Alma Mater Studiorum–Università di Bologna, Italy
{claudia.cevenini, enrico.denti, andrea.omicini, italo.cerno}@unibo.it

Abstract Large-scale *socio-technical systems* (STS) inextricably interconnect individual – e.g., the right to privacy –, social – e.g., the effectiveness of organisational processes –, and technology issues —e.g., the software engineering process. As a result, the design of the complex software infrastructure involves also non-technological aspects such as the legal ones—so that, e.g., law-abidingness can be ensured since the early stages of the software engineering process.
By focussing on *contact centres* (CC) as relevant examples of knowledge-intensive STS, we elaborate on the articulate aspects of *anonymisation*: there, individual and organisational needs clash, so that only an accurate balancing between legal and technical aspects could possibly ensure the system efficiency while preserving the individual right to privacy. We discuss first the overall legal framework, then the general theme of anonymisation in CC. Finally we overview the technical process developed in the context of the BISON project.

Keywords: socio-technical systems; contact centres; anonymisation; privacy

1 Introduction

Socio-technical systems (STS) are those systems where "the infrastructure is technology, but the overall system is personal and social, with all that implies" [12]. Large-scale STS [9] are nowadays typically characterised by a large number of participants and components, as well as by a huge amount of available data—recorded, produced, and used by the system activities.

Among the most relevant cases of large-scale STS are *contact centres* (CC)—in particular in Europe, where they involve nearly 1% of its active population. CC are clearly *knowledge-intensive* systems, since they typically produce a wealth of spoken data, which are mined either manually or by rudimentary technical means. Spoken data in large-scale European CC are often *multilingual* and involve *multiple countries*, meaning that both national and EU laws and regulations on personal data and privacy have to be taken into account. In particular, processing of personal data is only performed when necessary, and by prior obtaining the *data subject*'s consent for the specific processing purpose.

Typical technology issues of CC as STS involve (1) basic speech data mining technologies with multi-language capabilities, (2) business outcome mining from speech, and (3) CC support systems integrating both speech and business outcome mining in user-friendly way. Scaling up to big (i.e., massive) data processing clearly scales up also the privacy and data protection issues. Moreover, when industrial research is performed, a distinction has to be made between the research phase – when software and technologies are being developed and tested – and the subsequent market stage—when real customer data are processed. These are the very motivations behind this paper: that is, how complex legal issues at both national and international level can be dealt with while building a complex software infrastructure for CC—both in the development and in the subsequent business phases. So, first of all, this paper aims at investigating how complex software infrastructures for CC may be developed and marketed in the full respect of the data protection legal framework.

The legal analysis should thus necessarily complement and support the technical work since the very early stages, acting as an enabler rather than an obstacle, by providing the legal framework within which a CC system may be developed and used. The analysis should: identify and analyse the legal requirements of speech data processing systems; investigate the relevant legal framework; determine the impact of legal and ethical issues on the deployment of the CC infrastructure; examine how such a system should be designed and used so as to comply with the applicable legal framework, while identifying the barriers that could potentially affect its design and deployment; and, finally, keep an eye over the procedures for data collection, storage, protection, retention, and destruction, so that they comply with national and EU legislation.

In this paper we focus on *anonymisation* [7] as a fundamental tool to deal with the potential conflict between opposite rights and needs, especially in the research and development phase of a large-scale, knowledge intensive STS. Conceptually speaking, anonymising amounts at defining *which* and *how much* information should be removed for some data to be acceptedly considered as anonymous—i.e, not de-identifiable [1] with "normal" means; technically speaking, an effective anonymisation process needs to suitably balance the effort for anonymising data with *(a)* the value of the resulting data, and *(b)* the purpose for which they are collected and used. In fact, while in principle the total anonymisation of personal data would obviously address the users' desire for privacy, the full availability of (spoken) data is often essential for the organisation to fulfil its goals—and at least useful for the efficiency of its processes.

As a result, the need for a suitable compromise between law-abidingness (and privacy needs), on the one side, and system and process efficiency, on the other, is a relevant goal not just for the legal analysis, but for the whole engineering process that leads to the construction of the CC infrastructure, so that a potential conflict of interests becomes *composition* of interests, and the law-abidingness requirement can be exploited as a *success factor* instead of being perceived as a possible source of delays and overheads—an issue that goes well beyond the (noteworthy) case study discussed here.

In the remainder of this paper we first recall the main legal issues (Section 2), then perform a socio-legal-technical analysis aimed at identifying the most relevant principles (about data protection and processing, about security measures, and others) and the consequent technological requirements (Section 3) as a pre-requisite to frame and discuss in depth the anonymisation process—first in general terms (Section 4), then in the specific context of the BISON project (Section 5). There, the specific goal is to understand how to structure the anonymisation process during the industrial research phase, yet without compromising the quality of development and testing, which is based on the data used, allowing the resulting STS to eventually deal with the proper amount of data when it reaches the business operation phase.

2 Legal Framework

Data protection is a fundamental human right, recognised by Council of Europe Convention – Treaty 108 [11], the first legally binding international instrument for data protection, by the Treaty on the Functioning of the European Union [6], and by the Charter of Fundamental Rights of the European Union [2].

The legal framework at EU level is laid down by Directive 95/46/EC (Data Protection Directive, or *DPD* [4]). Brought into force by all EU Member States' national law, the DPD contains key principles for the fair and lawful processing of personal data, together with the technical and organisational security measures designed to guarantee that all personal data are safe from destruction, loss, alteration, unauthorised disclosure, or access, during the entire data processing period. Data processing requires even more care when it involves *large amounts* of personal and/or sensitive data: in particular, people should be given the possibility to manage the flow of data relating to them across massive, third-party analytical systems, so as to have a transparent view of how information data will be used, or sold.

The *data transfer from and outside the EU and cloud services* is therefore a particularly hot topic, since non-EU countries might provide an insufficient level of protection to personal data. This is why the flow of personal data is free between EU Member States, whereas the DPD sets restrictions on the export of personal data to third countries not ensuring an *adequate* level of data protection. Adequacy of data protection in a third country means that the main principles of data protection are effectively implemented in the national law of that country. Therefore, when there is no specific consent to the data transfer outside the EU given by the data subject, and when the level of data protection of the recipient's country is not deemed adequate, the data controller may be required – before exporting personal data – to contractually bind the recipient to set up enough security and organisational measures to grant adequate protection of the personal data—e.g., through standard contractual clauses, binding corporate rules.

2.1 Personal data and (de-)identification

Personal data consist of any information relating to a natural person, who can be identified, either directly or indirectly, by reference to one or more factors specific to his/her physical, physiological, mental, economic, cultural, or social identity. It should be noticed that if the link between an individual and his/her data never occurred, or, it is somehow broken and cannot be rebuilt in any way (as in the case of anonymised data), the DPD rules no longer apply: this is why anonymisation turns out to be a fundamental tool to simplify both the industrial research process and the processing system design and development—clearly, in as much as the data value is not compromised. With respect to this issue, it is worth recalling the Explanatory Report to Convention 108 [11], which states that

- "identifiable person" means a person who can be *easily* identified: it does not cover identification of persons by means of 'very sophisticated methods' (Article 2, § 28);
- 'the requirement appearing under litterae concerning the time-limits for the storage of data in their name-linked form does not mean that data should after some time be *irrevocably separated* from the name of the person to whom they relate, but only that it should not be possible to *link readily* the data and the identifiers' (Article 5, § 42).

As concerns the collection of personal data (the very first processing operation), the DPD sets out some basic definitions and principles for lawful processing.

First, the DPD identifies two distinct roles: the *data controller* and the *data processor*. The former is in charge of personal data processing and takes any related decision—e.g. selection of data to be processed, purposes and means of processing, technical and organisational security, etc. The latter, instead, is a legally separate entity that processes personal data on behalf of a controller, in force of a written agreement and following specific instructions. In other words, the controller processes data on its own behalf, while the processor always acts on behalf of a controller, from whom it derives its power and range of activity. For instance, a company acts as a controller in processing its own customers' data, whereas the CC entrusted with the same processing acts as a processor on behalf of the company.

Personal data must be obtained and processed lawfully, be collected for *explicit* and *legitimate* purposes, and *used accordingly*. The processing purposes must always be clearly declared, primarily to the data subject, who has to be specifically informed of all elements related to the processing, before the processing itself is started: the data subject has then to to provide his/her free, specific, and unambiguous *consent*. Any processing for undefined purposes is not law-abiding, and the consent given in such cases is not deemed valid. The same applies when the data subject is asked for just one consent in view of a plurality of purposes.

The data controller must not use the data collected for a given purpose to pursue a different one, also at a different time (i.e. after the declared purpose has been achieved): any further use of personal data for other purposes requires

an additional legal basis if the new purpose of processing is incompatible with the original one. Furthermore, the data collected must be *strictly consistent* with the declared purposes: it is unlawful to collect more data than necessary (a.k.a. *principle of necessity*). Any data transfer to third parties is also a new purpose, potentially requiring additional legal support.

Personal data must also be *relevant* and *not excessive* in relation to the purposes for which they are collected and processed: only the specific data which are actually necessary to achieve a given purpose may be collected and processed, any wider collection resulting in law infringements. Personal data must also be *accurate* and *up to date*: whenever pieces of information on a given data subject turn out to be wrong, or need to be changed, the personal data must be consequentially corrected.

From the timing viewpoint, data may be retained *only for the period needed* to achieve the specific purposes for which they are being processed: then, they should be erased. However, it is possible to continue the data processing beyond the originally declared purposes if personal data are anonymised—that is, they cannot be linked back to an individual in any way. So, the anonymisation process may be regarded to as the last authorised operation of the processing of personal data, before they cease being personal data to become simply "data".

2.2 Accountability and security measures

According to the *accountability* principle, data controllers have to implement adequate measures to promote and safeguard data protection in their processing activities. Controllers are responsible for the compliance of their processing operations with data protection law, and should be able to demonstrate compliance with data protection provisions at any time. They should also ensure that the practical measures implemented to comply with data protection principles are effective. In case of larger, more complex, or high-risk data processing, the effectiveness of the measures adopted should be verified regularly, through monitoring, internal and external audits, etc.

Technical and organisational *security measures* should be adopted to protect personal data, during all the processing period, against the risks related to the integrity and confidentiality of data, in particular where the processing involves the transmission of data over a network, and against all other unlawful forms of processing. The level of data security requested by the law is determined by different elements, such as the nature (sensitive/non-sensitive) of the collected data, the concrete availability in the market of adequate security measures at the current state of the art, and their cost—which should not be "disproportionate" with respect to the necessity.

It is worth pointing out, however, that security measures are not limited to technical remedies, but also include organisational rules and procedures that should be strictly observed by all the subjects involved in data processing. The overall quality of the security measures is the result of a case-by-case evaluation, which must be performed before starting new personal data processing operations, and then implemented and adapted, when needed, during the whole data

processing period, with regard to the technical solutions, and the human factor, too.

2.3 Big (Speech) Data

A CC infrastructure involves *speech recordings*, that is, processing biometric data (as in the case of the analysis of the tone, pitch, cadence, and frequency of a person's voice) for determining whether a person is who he/she declares to be.

From a data protection perspective, biometric technologies are closely linked to physical, physiological, behavioural, or even psychological characteristics of an individual—and some of them may be used to reveal sensitive data. Biometric data may also enable automated tracking, tracing, or profiling of persons: as such, their potential impact on privacy is quite relevant. Moreover, biometric data are by nature *irrevocable*: a breach concerning biometric data threatens the further safe use of biometrics as identifier, as well as the right to data protection of the concerned persons for whom there are no chances to mitigate the effects of the breach.

Therefore, the processing of biometric data is not only subject to the express consent of the data subject, but may also depend on the authorisation by Data Protection Authorities, and is submitted to strict rules on security measures that must be adopted to protect data: for instance, biometric information should always be stored in encrypted form; decryption keys should only be accessible on a need to know basis; the system should be designed in a way that allows the identity link to be revoked, either in order to renew it or to permanently delete it if the consent of the data subject is revoked; etc.

In this context, large-scale STS such as CC deal with big data because of the huge amount of data they collect, even though from a limited number of sources. Two main issues are worth highlighting:

- Big data analytics can involve the *repurposing* of personal data. If an organisation has collected personal data for one purpose and then decides to start analysing it for another one (or to make it available for others to do so), data subjects need to be informed of this novelty, and a new, specific consent is usually needed. This is particularly important if the organisation is planning to use the data for a purpose that is not apparent to the individuals because it is not obviously connected with their use of a service.
- Big data may intrinsically contrast with the principle of data minimisation and relevancy: the challenge for organisations is to focus on what they expect to learn, or, to be able to do by processing big data before the beginning of processing operations, thus verifying that these serve the purpose(s) they are to be collected for, and, at the same time, that they are relevant and not excessive in relation to such aim(s).

3 Socio-Legal-Technical Analysis

The current legal framework foresees a set of essential principles that should inspire the design and development of any law-abiding information system pro-

cessing personal data. While some of such principles directly derive from the DPD – namely, from the "Principles relating to data quality" –, others concern the security measures that should be adopted, particularly with reference to the "Security of processing". These principles are further strengthened and detailed in the "General Data Protection Regulation" (GDPR) [8].

3.1 Relevant principles

Relevant principles can be conceptually organised in three major categories, which are shortly detailed in the following:

(a) principles about data processing
(b) principles about security measures
(c) other relevant principles

Principles about data processing

1. *Principle of lawfulness and fairness*: any processing of personal data must be lawful and fair to the individuals concerned.
2. *Principle of relevance and non-excessive use*: personal data must be adequate, relevant and not excessive in relation to the purposes for which they are collected and/or further processed.
3. *Principle of purpose*: personal data must be collected for specified, explicit and legitimate purposes and not further processed in a way incompatible with those purposes.
4. *Principle of accuracy*: data must be accurate and, where necessary, kept up to date; every reasonable step must be taken to ensure that data which are inaccurate or incomplete, having regard to the purposes for which they were collected or for which they are further processed, are erased or rectified.
5. *Principle of data retention*: data must also be kept in a form which permits identification of data subjects for no longer than is necessary for the purposes for which the data were collected or for which they are further processed; Member States shall lay down appropriate safeguards for personal data stored for longer periods for historical, statistical or scientific use.

Principles about security measures

1. *Principle of privacy by design*: the protection of the rights and freedoms of data subjects with regard to the processing of personal data requires appropriate technical and organisational measures, at the time of the design of the processing system as well as at the time of the processing itself, particularly in order to maintain security and thereby to prevent any unauthorised processing; these measures must ensure an appropriate level of security, taking into account the state of the art and the costs of their implementation in relation to the risks inherent in the processing and the nature of the data to be protected.

2. *Principle of appropriateness of the security measures*: the controller must implement appropriate technical and organisational measures to protect personal data against accidental or unlawful destruction or accidental loss, alteration, unauthorised disclosure or access, in particular where the processing involves the transmission of data over a network, and against all other unlawful forms of processing. Having regard to the state of the art and the cost of their implementation, such measures shall ensure a level of security appropriate to the risks represented by the processing and the nature of the data to be protected.
3. *Principle of privacy by default*: the controller shall implement mechanisms for ensuring that, by default, only those personal data are processed which are necessary for each specific purpose of the processing and are especially not collected or retained beyond the minimum necessary to achieve those purposes, both in terms of the amount of the data and the time of their storage. In particular, those mechanisms shall ensure that by default personal data are not made accessible to an indefinite number of individuals.

Other relevant principles

1. *Principle of least privilege*: in a given abstraction layer of a computing environment, every module (process, user, program) must be able to access only the information and resources that are necessary for its legitimate purpose.
2. *Principle of intentionality in performing any critical action*: examples include granting access to a wider set of users, selecting lower security settings, exporting data, reducing the number of anonymised items, etc.

3.2 Consequent technological requirements

The above principles translate into actual system requirements ranging from the system configuration to the user management, the way data are processed, and the security measures in general. Abstracting from any specific technical solution, the following issues can be identified and should be accounted for:

1. Different user groups and different users, with different privileges and access rights, so that each user is granted only the minimum set of rights that is necessary for his/her task, coupling maximum flexibility with maximum security; this asks for role-based authentication model, fine-grained set of user rights, adequate authentication mechanisms.
2. Default user profiles with the minimum set of rights, so that any addition of user rights giving access to data is always intentional.
3. Default anonymisation configuration corresponding to the maximum level of anonymisation, so that any custom configuration implying a decrease of anonymisation level is always explicitly authorised, and therefore intentional and security-checked before proceeding.
4. Adequate support of detailed customised anonymisation levels, so that the system settings can be fine-tuned to the specific customer necessity—and no further.

5. Multiple levels of security, with proper warnings, whenever the default (i.e., maximum security) settings are lowered for any reason—e.g. during the custom configuration by authorised and suitably authenticated personnel.
6. Clear identification of the use-case scenarios when (authorised) personnel is allowed to access personal data—that is, non-anonymised copies of recordings, lawfully stored for authorised processing.
7. Severe restrictions on, or even denial of, transfer of non-anonymised data.
8. Immediate removal/wiping of non-anonymised copies of data recordings that may have been made during the processing, if required by the processing itself, as soon as their presence is no longer required.

4 Anonymisation Process

As detailed above, the legal framework allows personal data to be processed only to the extent they are needed to achieve specific purposes: whenever identifying data are not necessary, only anonymous data should be used.

As far as legal requirements for anonymisation are concerned, the DPD does not apply to data made anonymous in such a way that the data subject is *no longer identifiable*: it does not set any prescriptive standard, nor does it describe the de-identification process—just its outcome, which is a *reasonably-impossible* re-identification. The concrete application of such a general principle, however, is not easy: the main problem is to create a truly anonymous dataset, while retaining at the same time all the data required for a specific (organisational) task. On the other side, irreversibly-preventing identification requires data controllers to consider all the means which may 'likely reasonably' be used for identification, either by the controller or by a third party.

The Directive 2002/58/EC (e-Privacy Directive) [5] also imposes anonymisation in certain cases. For instance, as far as subscribers' traffic data are concerned, processed within electronic communications networks to establish connections and to transmit data, it foresees that, when used for marketing communications services or value added services, personal data should be erased or made anonymous after the provision of the service. Besides, it imposes the providers of a public communication network or a publicly-available electronic communication service to erase or anonymise data no longer needed to transmit a communication.

The *Article 29 Working Party – Opinion on Anonymisation Techniques* (Art. 29 WP henceforth) [3] is an important reference for compliance in anonymisation issues: it describes the main techniques used to anonymise personal data, and explains their principles, strengths and weaknesses, possible mistakes, and failures. The criteria on which Art. 29 WP grounds its opinion on robustness focus on the possibility of:

- singling out an individual;
- linking records relating to an individual;
- inferring information concerning an individual.

Assuming that personal data have been collected and processed in compliance with applicable legislation, Art. 29 WP on the one hand considers anonymisation as 'further processing', which generally speaking needs to comply with the compatibility assessment—e.g. data shall not be further processed for purposes incompatible with those specified at the moment of their first collection; on the other, however, it promotes the idea that anonymisation should be seen as further processing *compatible with the original purposes*, upon condition that the anonymisation process *reliably produces anonymised information*.

Again, a balance between different needs has to be found: although on the one hand removing directly the identifying elements in itself may not be enough to ensure the impossibility of re-identification, the principle set by Convention 108, on the other, states that identification *with very sophisticated methods* may not be relevant, as it should not be possible to identify a person *readily*.

Additional measures may often be needed to prevent identification, depending on the context and purposes of the processing for which the anonymised data will be used. In its Opinion 03/2013 on purpose limitation, Art. 29 WP notes that

> Anonymisation is increasingly difficult to achieve with the advance of modern computer technology and the ubiquitous availability of information. Full anonymisation would also require, for instance, that any reasonable possibility of establishing a link with data from other sources with a view to re-identification be excluded. However, re-identification of individuals is an increasingly common and present threat. In practice, there is a very significant grey area, where a data controller may believe a dataset is anonymised, but a motivated third party will still be able to identify at least some of the individuals from the information released. Addressing and regularly revisiting the risk of re-identification, including identifying residual risks, therefore remains an important element of any solid approach in this area.

Adequate safeguards, whose strength should be proportionate to the adverse impact to the data subject of a possible re-identification, should also be considered if necessary—such as encryption, restrictions of access, etc.

More generally speaking, it is hard to assess in advance whether re-identification may or may not be possible, since it depends on how readily and how directly the link between the data and the identifiers is structured. Data controllers, for instance, should consider the *concrete means* that would be necessary to reverse the anonymisation process, their cost and know-how needed, as well as the likelihood and severity of implementing such means. They should conduct a data protection impact assessment to decide what data may be made available for reuse and at what level of anonymisation and aggregation: ideally, an impact assessment should be completed before disclosing information and making it available for reuse. Whenever controllers release anonymised datasets for use, the risk assessment should include re-identifiability tests (e.g. penetration tests). Finally, data controllers should keep into account that the risk of re-identification changes over time, with the evolution of technology: once-rare

and sophisticated analytics techniques can quickly become commonly available, possibly even at low or no cost, or, new evidence could reveal accessible techniques for re-identification—like, e.g., in [13]. Thus, the data controller policies should be periodically reviewed in consideration of current and possible future threats.

Data processors, on the other hand, may process anonymised datasets communicated to them by data controllers: if they are not able to either directly or indirectly identify the data subjects of the original dataset, they will be acting lawfully with no need to consider data protection requirements. Still, before deciding if and how to use the anonymous data received by the data controller for their own purposes, they should evaluate the anonymisation techniques adopted by the data controller, because data processors may be held liable for consequences derived from their own data processing. Thus, if there is a risk of identification of the data subjects, the processing should be performed in compliance with the data protection law.

From the technical viewpoint, two aspects need be stressed: *(i)* different anonymisation techniques may be used, which may imply different levels of risk; *(ii)* anonymisation is necessarily defined with respect to some *threshold* about the easiness or probability of singling out, linking, or inferring an individual in a dataset—that is, when some data are considered sufficiently *de-identified*.

Art. 29 WP – Opinion 05/2014 outlines the risks of singling out, linking, or inferring, with respect to some of the most used techniques (e.g. pseudonymisation, noise addition, substitution, etc.); in [1], a different approach, rooted in the USA, is discussed, based on the number k of *quasi-quantifiers* (i.e., identity-revealing traits) that must be cancelled in a dataset so as to ensure that the probability of re-identification is below a given threshold: this leads to the concept of *k-anonymisation* (there, a value of 3 is considered the minimum, with 5 being reasonably safe for the purpose). Also, this fits perfectly the other observation pointed out by Art. 29 WP, i.e., that anonymisation can protect privacy and personal data only if anonymisation techniques are engineered and applied properly: context and objectives of the process must be clearly set to achieve the desired/required anonymisation level.

5 The Anonymisation Process in BISON

The issues discussed in this paper have been investigated in the context of the BISON project [10], aimed at developing an innovative tool for CC processing big speech data.

As it is common in industrial research, a fundamental distinction has to be pointed out between the *research phase* – when software and technologies are being developed and tested, but are not yet in actual production – and the subsequent, foreseeable *business phase* – when they actually deal with real customers' data.

Here, anonymisation is seen as the fundamental tool to set the industrial research phase free from the complex requirements imposed by the Data Pro-

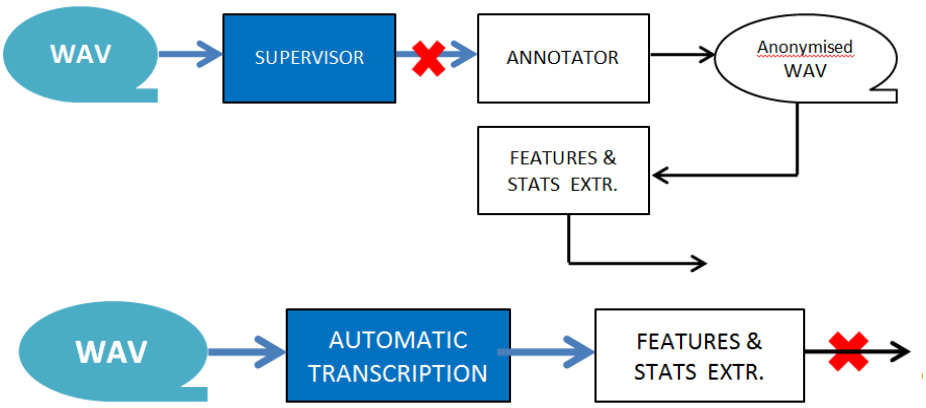

Figure 1. Anonymisation during the Start-up stage and Research stage in BISON

tection rules, given that the DPD does not apply to anonymised data. At the same time, in the business phase that will follow the research project, the tool will have to deal with real user data, in compliance with applicable laws.

5.1 General overview

In the first stage of the BISON research, data flow as in Figure 1 (top): the anonymisation process is performed mostly with manual procedures, both because of the limited data size and because of the initial lack of automatic tools. The starting point is the audio file (WAV) of the call—which contains personal data. The call is examined by a *supervisor*, who is the only person authorised to access personal data, with the specific purpose of intercepting any personal data in the audio call and manually removing them by 'silencing' (that is, by physically overwriting the relevant words with silence or some suitable "beep"). The result is a new audio file that contains no personal data: therefore, from this point in the data flow (red cross in the diagram), data can be considered as fully anonymised. The anonymised call can then be passed to the *annotator*, who is in charge of tagging the file with keywords, according to pre-defined technical specifications: the resulting annotated file is then further processed to extract statistics (features), which embed enough information to train the audio recognisers while making it impossible to reconstruct the original waveform and to trace back to the original personal data.

In the second stage of the BISON research, huge amounts of speech data need to be processed, which makes the manual annotation of personal data by the supervisor unfeasible: therefore, automatic transcription — for all the supported languages — has to be put in place. This change affects the stage where anonymisation takes place, since anonymisation is now performed on the original audio file (containing personal data) instead of on a manually pre-silenced

audio file; it may also somehow reduce the reliability of the process, since no automatic transcriber can be considered 100% effective in identifying terms and items related to any possible personal data: therefore, the "anonymised" file in this case may still contain some (hopefully and typically, a small amount of) personal data.

Of course, any effort should be made to reduce these errors to the minimum: thus, the automatic anonymiser should be designed and trained with the greatest possible care, and tested according to the best available practices. The subsequent feature extraction helps to deal with this issue, because the extracted statistics make it (mostly) impossible to reconstruct the original audio file. This is why the red cross in Figure 1 (bottom) is conceptually placed only after the feature extraction step, instead of following the automatic transcription—although most of personal data are actually suppressed earlier in the chain. Features are finally exploited to feed the language recognisers with big data in many different languages (possibly under the requirement of signing a Non-Disclosure Agreement): in any case, data anonymisation takes place prior to the annotation of the recordings, which guarantees full anonymisation afterwards.

In a farther perspective, when the final system will eventually operate in the business context, the basic difference will concern *where* the extraction of statistics will be performed—namely, inside each CC, by the CC itself: so, no personal data will ever be delivered outside, and any processing will occur only provided that the appropriate consent, for the specific purpose, has been given, and in compliance with all applicable laws and regulations.

5.2 Technological requirements

Despite the basic assumption of relying mostly on automatic anonymisation, some manual adjustments might still be necessary in the development and configuration phases—and possibly even at runtime, so as to capture any residual data that might have survived the previous checks. For this reason, automatic technologies should be coupled with an interactive tool, enabling the fine-tuning and (possibly live) control of the anonymisation process.

Such a tool should obviously adhere to strict security requirements: users' roles, rights, and restrictions should be tuneable on a fine-grain basis, and be further detailed case-by-case based both on the actual needs and the applicable national legal framework. Moreover, on-the-fly anonymisation should be available to deal with the case that some unexpected personal data are heard by the CC agent in charge of the call, requiring real-time anonymisation to be triggered.

5.3 Customisation and future-proofing

In the final state of the system (ready-to-market), users will need to be enabled to anonymise personal data whenever not needed for the specific purposes of the processing—and they should be able to do so in a *highly customisable* way. Customisation should be based on the specific CC requirements: for instance, it should be possible to enable anonymisation at different times (during/after the

call), or based on the occurrence of specific situations, or as a feedback from speech analytics or data mining on text, etc.

A key challenge from this viewpoint is also to make anonymisation *future-proof* both with respect to a continuously-evolving legal scenario, as well as to the technology improvement, evolving even faster.

6 Conclusion

It is nowadays taken as understood that the practices of contemporary software engineering have to be extended to include non-computational issues such as normative, organisational, and societal aspects. This holds in particular for large-scale socio-technical systems: for instance, the law-abidingness of complex software systems including both human and software agents is quite an intricate issue, to be faced in the requirement stage of any reliable software engineering process.

In this paper we specifically address the problem of anonymisation of speech data in the case of contact centres, discussing the need for an accurate balancing between legal and technical aspects in order to ensure the system efficiency while preserving the individual right to privacy, and showing how the legal framework can actually translate into requirements for the software engineering process. By discussing the BISON approach, we show how the anonymisation process can be structured during the industrial research phase in order to make it possible for the resulting system to eventually deal with the amount of data actually required once it reaches the business operation phase.

Acknowledgements. This work has been supported by the EU-H2020-ICT-2014 Innovation action BISON – BIg Speech data analytics for cONtact centres (Grant Agreement no. 645323). Authors would like to thank all the partners of the BISON project for their invaluable cooperation in the development of the approach illustrated in the paper. Authors would also like to thank Dr. Silvia Bisi for her contribution to the project research.

References

1. Angiuli, O., Blitzstein, J., Waldo, J.: How to de-identify your data. Communications of the ACM 58(12), 48–55 (Nov 2015)
2. Charter of Fundamental Rights of the European Union. Official Journal of the European Communities 43(C 364), 1–22 (18 Jan 2000)
3. Article 29 Data Protection Working Party – Opinion 05/2014 on anonymisation techniques. `http://ec.europa.eu/justice/data-protection/article-29/` (18 Apr 2014), 0829/14/EN WP216
4. 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 38(L 281), 31–50 (23 Nov 1995)

5. Directive 2002/58/EC of the European Parliament and of the Council of 12 July 2002 concerning the processing of personal data and the protection of privacy in the electronic communications sector (Directive on privacy and electronic communications). Official Journal of the European Communities 45(L 201), 37–47 (31 Jul 2002)
6. Consolidated versions of the Treaty on European Union and the Treaty on the Functioning of the European Union. Official Journal of the European Communities 55(C 326), 1–390 (26 Oct 2012)
7. Fung, B.C.M., Wang, K., Chen, R., Yu, P.S.: Privacy-preserving data publishing: A survey of recent developments. ACM Computing Surveys 42(4), 14:1–14:53 (Jun 2010)
8. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) (text with EEA relevance). Official Journal of the European Communities 59(L 119), 1–88 (4 May 2016)
9. Omicini, A., Zambonelli, F.: Coordination of large-scale socio-technical systems: Challenges and research directions. In: Di Napoli, C., Rossi, S., Staffa, M. (eds.) WOA 2015 – From Objects to Agents. CEUR Workshop Proceedings, vol. 1382, pp. 76–79. Sun SITE Central Europe, RWTH Aachen University, Napoli, Italy (17–19 Jun 2015), proceedings of the 16th Workshop "From Objects to Agents"
10. The BISON Project: Home page. `http://bison-project.eu/`
11. Convention for the protection of individuals with regard to automatic processing of personal data. `http://conventions.coe.int/treaty/en/Treaties/Html/108.htm` (28 Jan 1981)
12. Whitworth, B.: Social-technical systems. In: Ghaou, C. (ed.) Encyclopedia of Human Computer Interaction, pp. 533–541. IGI Global (2006)
13. Zang, H., Bolot, J.: Anonymization of location data does not work: A large-scale measurement study. In: 17th Annual International Conference on Mobile Computing and Networking (MobiCom 2011). pp. 145–156. ACM, New York, NY, USA (2011)