Call identifier: H2020-ICT-2016 - **Grant agreement no**: 732907

Topic: ICT-18-2016 - Big data PPP: privacy-preserving big data technologies

Deliverable 3.2

Ethical Review materials and engagement process definition

Due date of delivery: 31st January, 2018

Actual submission date: 13th February, 2018

Start of the project: 1st November 2016
Ending Date: 31st October 2019

Partner responsible for this deliverable: LYNKEUS

Version: 0.8

Dissemination Level: Confidential Document Classification

Title	Dynamic Consent Ethical Review
Deliverable	D3.2
Reporting Period	M1 - M18
Authors	Mirko De Maldè, Ludovica Durst, Noel Catterall
Work Package	WP3
Security	PU
Nature	RE
Keyword(s)	

Document History

Name	**Remark**	**Version**	**Date**
Noel Catterall	Partial Preliminary Draft of section 3	0.1	September 13th 2017
Ludovica Durst	Additional text on dynamic consent	0.2	November 22nd, 2017
Mirko De Maldè	Additional text on smart contracts	0.3	December 18th, 2017
Mirko De Maldè	First overall draft	0.4	January 6th, 2918
Davide Zaccagnini	Review and comments	0.4	January 10th, 2018
Mirko De Maldè	Second overall draft	0.5	January 15th, 2018
Ludovica Durst	New revision	0.6	January 26th, 2018
Mirko De Maldè	Third Overall draft, taking into account gnúbila team review	0.7	January 29th, 2018
Mirko De Maldè	Final Draft	0.8	February 13th, 2018

List of Contributors

Name	**Affiliation**
Noel Catterall	HWC
Ludovica Durst	LYNKEUS

Mirko De Maldè	LYNKEUS

List of reviewers

Name	**Affiliation**
Davide Zaccagnini	LYNKEUS
Anna Rizzo	LYNKEUS
Mirko Koscina	Gnúbila
Alexandre Flament	Gnúbila
Federico Sartore	P&A
Lorenzo Cristofaro	P&A
Edwin Morley-Fletcher	Lynkeus

TABLE OF CONTENTS

1. PURPOSE OF THE DOCUMENT

The present deliverable is aimed at providing a general evaluation framework for the management of patient consent in MHMD, and its implementation as a smart contract. Such framework will serve as the basis for submitting in due course the needed requests, to MHMD partner hospitals' Ethics Committees, in view of being authorised to deploy dynamic consent applications within their patient populations.

If successful, the outcome of this activity will represent a major result of the project, as it will deliver one of the first ethically compliant (and approved) blockchain architectures specifically conceived for health data management and exchange.

The scope of this deliverable has been updated to take into account the changes in task T3.2, as indicated in the recently submitted Amendment, in particular to reflect the new strategy for implementing dynamic consent through smart contracts and leverage the blockchain to make consent management process (from the provision of consent to the ex-post verification of existence of such consent) transparent, semi-automatic and tamper proof.[1]

The change in strategy was deemed necessary to provide the Ethics Committees not only with the supporting documentation needed for approving a novel mechanism for expressing and managing consent (becoming fully dynamic through smart contracts), but also to encode the conceptual consent framework on which to obtain a comprehensive and robust ethical approval – incorporating the ways in which consent is implemented on the blockchain in code, shared with external parties, and automatically self-enacted and possibly modified by the patient.

With this document, an iterative process will be activated, using feedback gathered from Ethics Committees in order to better specify the entire process and corresponding IT architecture. The subsequent phases of the iteration are provided in section 4.4.5 (MHMD approach).

Beside presenting the issue of how to translate dynamic consent into smart contracts, this document will also provide a preliminary recognition and analysis of the technological approaches currently followed for smart contracts implementation, thus presenting a general development framework for applying the dynamic consent concept though smart contracts within MHMD.

[1] The filing of this deliverable (and of deliverable D.2.2) was slightly delayed in order to take into account, on the one hand, the provisions laid down by the Article 29 Working Party in its draft *Guidelines on consent* (WP259) and *Guidelines on transparency* (WP260), whose public consultation was closed only at the end of January 2018 and, on the other hand, some recent developments regarding the status of some Member States' national law implementing the GDPR.

1.1.USING SYNTHETIC DATASETS

To provide Ethics Committees with the intended conceptual and technological framework, MHMD needs to test the entire blockchain architecture and the relevant smart contracts execution process. This activity is planned to be completed in its first form (i.e., basic blockchain architecture and user interfaces for data exchange) within the first six months of 2018. To do this, though, actual datasets would be needed, while an ethical approval is needed to make datasets available. In view of avoiding being stuck in a sort of "catch 22" situation, an initial usage of synthetic datasets has been adopted in MHMD.

In particular, partner QMUL (at Barts' Hospital), has started to generate a synthetic cardiology data set belonging to fictitious individuals, based on aggregate statistics of a population of 100,000 patients. These datasets have spurious correlations added to reflect the impact of cardiovascular risk factors on cardiovascular health.

The datasets contain fake names, addresses, DOB, DOD, episode visits, anthropometry (e.g. weights, heights, BMI, BSA, etc.) and cardiac function parameters, etc. Examples of data types/sources targeted for early inclusion include Myocardial Infarction, cath lab data, demographics, CT images and radiology reports, MRI images and radiology reports, pacemaker data, echocardiography images and radiology reports, cardiac surgery reports, data from the chest pain clinic and pathology data.

It is worth mentioning that recent research[2] indicates as data science state-of-the art approach the provision of appropriately generated synthetic datasets aimed at replacing original data, given the fact that no significant difference can be spotted in the outcome of research based on the usage of synthetic data as opposed to real data.

This dataset, which has been already partially delivered by Barts, will allow MHMD to evaluate how standardized ontologies can be mapped onto such a (typical) data export format, to assess procedural requirements, such as loading and processing time, CPU requirements on multi-site computing, as well as

[2] Patki, N., Wedge, R., & Veeramachaneni, K. (2016, October). The Synthetic Data Vault. In *Data Science and Advanced Analytics (DSAA), 2016 IEEE International Conference on* (pp. 399-410). IEEE. A further example will soon be provided by SIMULACRUM, where simulated data with identical structure to real cancer registry data, including missing values and other artefacts, will be made public by Q1 2018, becoming suitable for study design and feasibility analysis. It will be possible to run identical queries on the real and fake data, making use of an expedited data release process: query results may be non-disclosive, despite relying on identifiable data. The latest iteration contains complex linked tables, with patients having multiple tumours, and detailed chemotherapy treatment data. (Dr Brian Shand, National Cancer Registration and Analysis Service, Public Health England, *Does one size fit all? Challenges of anonymising real world data*, presentation at the EMA Workshop on "Data Anonymisation_A Key Enabler for Clinical data Sharing", London, 30 Nov.-1 Dec. 2017).

evaluating the impact of pseudonymisation/obfuscation/aggregation techniques on a range of dimensional statistical measures.

Last but not least, the dataset will be used for the hacking challenges and penetration testing, while real data will need to be collected for reidentification tests, as soon as they are available and duly consented.

2. MHMD DATASETS CLASSIFICATION, AND CONSENT CONSIDERATIONS IN THE GDPR FRAMEWORK

Within MHMD, different datasets from several sources will be collected and used. Depending on the source and privacy level, different consent approaches will be adopted.

2.1. DATASETS USED IN MHMD

As more clearly indicated in the latest formulation of the Description of Action, MHMD will use both legacy data (coming from previous EU-funded projects such as MD-Paedigree and CARDIOPROOF, and from the hospital routine data collection process, after applying the most appropriate methods for ensuring data privacy and security, e.g., anonymisation, encryption, etc.), and data coming from individual patients enrolled during the project.

This distinction will also be valid when the system is fully deployed and available for external hospitals, as the scenario will envision two data types: the one that the hospital has already collected in the past (together with the needed consent) and is willing to share using the MHMD platform, and the data that will be "natively" shared through MHMD, meaning that the patient will be included in the loop, through the dynamic consent application, thus being able to participate in the decision regarding his/her data access policy definition, while also monitoring the access to his/her data, once shared. Besides these scenarios, it shall be also added the one that sees the individual patient/citizen joining the MHMD platform autonomously (e.g., through the app, for sharing health-related datasets coming from wearable devices or MIoT).

The aim of MHMD is providing hospitals with a tool for seamlessly managing permission for making de-identified (anonymised or pseudonymised) data accessible to external institutions (which is also a mandatory request for EU funded projects), while preserving privacy and security of datasets and patients (also by never allow identifiable data to leave hospitals secure internal IT environment).

MHMD will implement this permission system, leveraging specific Data access smart contracts for embedding preferences and permissions, regulating access to specific datasets and their acceptable usage, also including (when relevant) ethical clearance, data provenance, acknowledgment guidelines. At the same time, this permission system will also allow hospitals to manage more effectively the consent associated with each dataset and provide an advanced tool to interact with patients after having obtained the needed consent. In fact, for legacy and clinical routine data, MHMD will start from the assumption that the needed consent has been duly obtained by the hospital also for further research purposes, and thus that the hospital is lawfully sharing the associated pseudonymised dataset on the MHMD platform, also to third parties involved in consented research purposes. Should, however, the consent be referred exclusively to research activities carried out by the hospital-controller itself, then the data, before being shared with third parties, would need to be anonymized. At the same time, hospitals might want to digitally operationalise the acquired consent to manage their permission system more effectively, attaching the original consent (with the specification about admitted data usage, specific purposes, possible expiration time, and other patient-driven data access preferences) to the Data access smart contract. This system, that hospitals could decide of adopting *ex post* on the already existing consented datasets, might also be adopted routinely for managing future consent for prospective patients.

For already consented data, this system will thus allow to manage consent digitally, making preferences interacting with the MHMD data platform through relevant smart contracts. Most importantly, it will make possible for hospitals to contact patients back, in case *ad hoc* consent is needed for specific purposes not explicitly consented at the time, in compliance with the General Data Protection Regulation no. 2016/679 – "GDPR" (as explained below – Section 5.1).

For this permission consent system, we aim to implement the following scheme (developed within D3.1), representing an interface for a hospital IT director/data officer of a hospital.

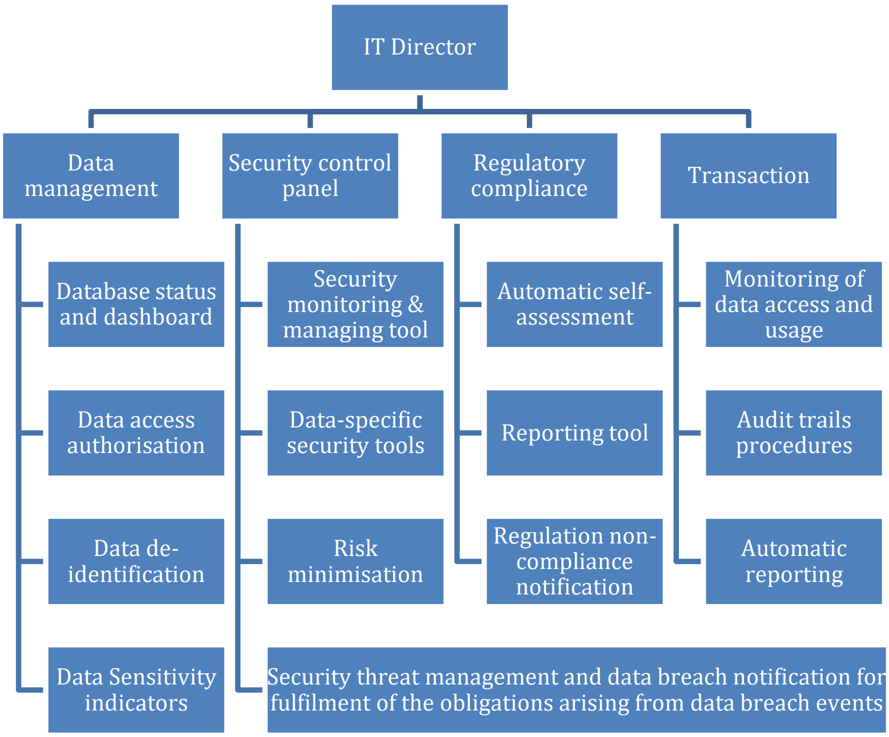
IT Director
Data management
Security control panel
Regulatory compliance
Transaction
Database status and dashboard
Data access authorisation
Data de-identification
Data Sensitivity indicators
Security monitoring & managing tool
Data-specific security tools
Risk minimisation
Security threat management and data breach notification for fulfilment of the obligations arising from data breach events
Automatic self-assessment
Reporting tool
Regulation non-compliance notification
Monitoring of data access and usage
Audit trails procedures
Automatic reporting

Contextually, a preliminary list of permission options relating to pseudonymised data has been conceived as follows:

- Pseudonymised datasets available for research
    - For a specific disease (e.g. only for diabetes)
    - Not for: (e.g. exclusion)
    - For a given period of time (available until xx/xx/xxxx)
    - Secondary research usage consented (Y/N)
        - for research carried out by the original hospital
        - for research carried out by a third party
    - Specific research type/research centre category (private-public)
- Pseudonymised datasets available for Clinical Trials
    - Virtual cohort composition (Y/N)
    - Consent to contact the patient (Y/N)
- Pseudonymised datasets available to third parties for industrial usage:
    - For a specific disease: specify

  - For specific category of tools (drugs/device)
  - Not for: (e.g. exclusion)
  - For a given period of time (available until xx/xx/xxxx)
- Profiling (Allow profiling: Y/N)
- Statistical analysis (Allow statistical analysis on the data Y/N)

The resulting permission system interface design is illustrated below.

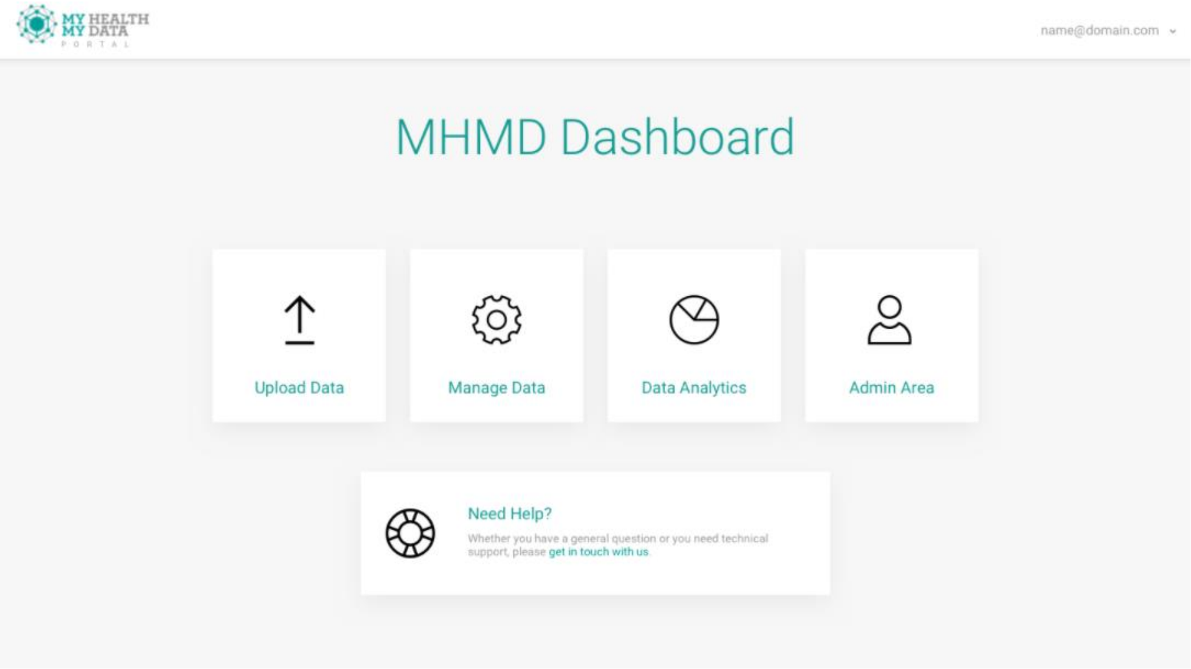

Whether you have a general question or you need technical support, please get in touch with us

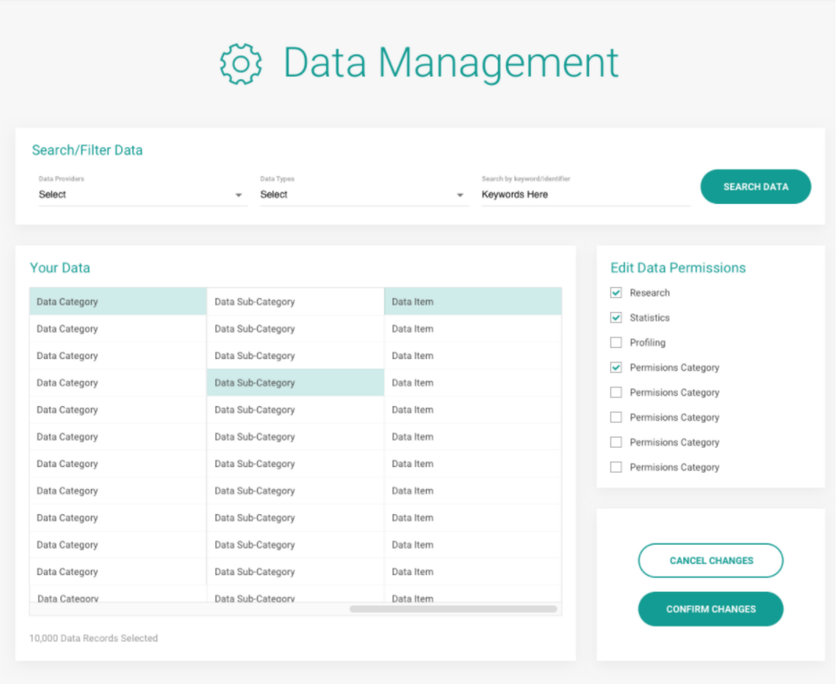

Data Providers: Select
Data Types: Select
Search by keyword/identifier: Keywords Here

Data Category	Data Sub-Category	Data Item
Data Category	Data Sub-Category	Data Item
Data Category	Data Sub-Category	Data Item
Data Category	Data Sub-Category	Data Item
Data Category	Data Sub-Category	Data Item
Data Category	Data Sub-Category	Data Item
Data Category	Data Sub-Category	Data Item
Data Category	Data Sub-Category	Data Item
Data Category	Data Sub-Category	Data Item
Data Category	Data Sub-Category	Data Item
Data Category	Data Sub-Category	Data Item
Data Category	Data Sub-Category	Data Item

10,000 Data Records Selected

☑ Research
☑ Statistics
□ Profiling
☑ Permisions Category
□ Permisions Category
□ Permisions Category
□ Permisions Category
□ Permisions Category

For new data, all these features would be available for both hospitals and patients, particularly allowing the latter to directly interact with their consent preferences, updating them at any time as per the dynamic consent approach. In case of conflict, it is generally necessary to refer to the principle of chronological sequence of consents, meaning that – as a general rule – "the last is the best", though some exceptions may apply on a case by case basis.

The data subjects directly providing their data will be able to exert their dynamic consent by making use of the dedicated MHMD Mobile application, specifically conceived to express and manage consent, automatically translating sets of preferences in smart contracts, regulating automatically data access and allowed usage.

This mobile application will provide an overview of the available patient datasets, of their sharing/access preferences, and a smart contract management dashboard, as illustrated in the figure below.

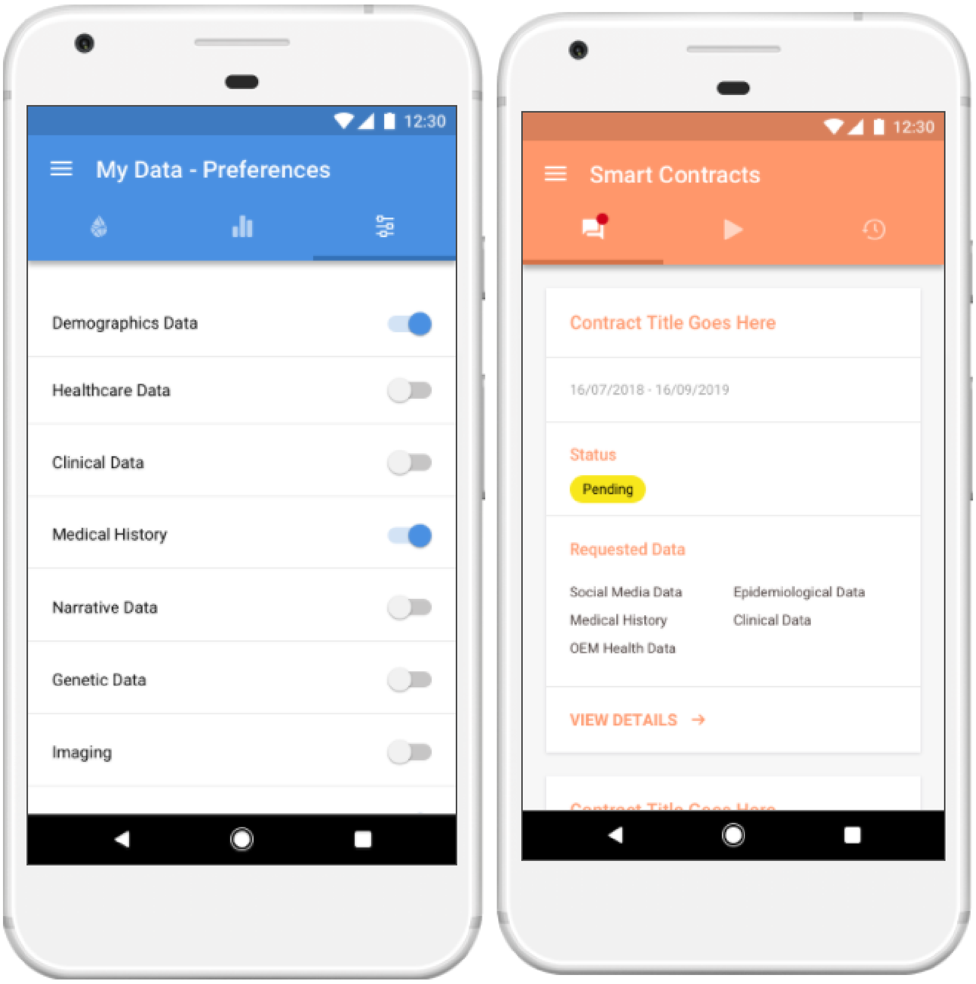

Demographics Data ●
Healthcare Data ○
Clinical Data ○
Medical History ●
Narrative Data ○
Genetic Data ○
Imaging ○

Contract Title Goes Here
16/07/2018 - 16/09/2019
Status
Pending
Requested Data
Social Media Data
Medical History
OEM Health Data
Epidemiological Data
Clinical Data

Both the hospital permission system and the patient mobile app are currently under development, and the first design (including key features) will be presented within the forthcoming deliverable D3.3 (due by M16).

2.1.1. Summary of data and associated consent used in MHMD

Summarising what explained above (and in line with section “Hospital Data Sharing Requirements” of the DoA), MHMD will make it possible to manage two different types of data, with associated consent needs:

- de-identified retrospective data (“pre-MHMD”, i.e., data already collected and available in the hospitals databases, for which each hospital already obtained consent, and which have been pseudonymised to allow re-use for research);
- prospective data with consent from individuals (“MHMD-native”).

For the first kind of data, two lawfulness conditions are needed to proceed with the processing:

1) patients have been clearly informed about the processing of their data for research purposes;
2) their consent has been acquired for a specific purpose.

If both conditions are met, pseudonymisation is considered enough. If only one of these conditions is not met, the hospital will need to have recourse to anonymisation for sharing the data.

For the MHMD-native data, the chosen de-identification method will depend on the classification of the data and on the relevant selection of the most appropriate privacy preserving technology, though generally all data will still belong to the pseudonymisation category, even if encrypted, and will imply making use of the consent management and security measures implemented by MHMD .

The following table, which is the result of a number of technical meetings where various use case scenarios have been explored, summarises data sources, the level of data security, and the relevant applicable consent policy.

Source	**Datasets**	**Consent**
Hospitals	Personal (non-deidentified) datasets collected in the clinical practice (all kind of data)	The consent is routinely obtained by the hospital for the specific treatment needed by the patient. Such raw data will never be shared outside the hospital.
Hospitals + privacy preserving processing	Personal pseudo-anonymised datasets collected in the clinical practice (all kind of data)	The needed consent is obtained by the hospital for the patient’s specific treatment and/or for further research purposes, in compliance with recital 33 of the GDPR. For secondary use, dedicated options and specific information shall be provided to the patient.
Hospitals + privacy preserving processing	Anonymised datasets derived from personal data collected in the clinical practice (all kind of data)	No consent is needed for sharing anonymised data. Hospitals do not need to be involved in the technical aspects of anonymisation: a dedicated tool will be provided by Gnúbila.

Hospital + generation of synthetic data	Synthetic data, statistically developed from real data	Synthetic data are by definition comparable to anonymised data. No consent is needed. A first dataset based on aggregate statistics of a population of 100,000 patients is provided by QMUL.
Citizens/patients	Personal data collected through wearable devices, mIoTs, etc.	Consent is needed and has to be collected directly by the relevant citizens'/patients' mobile application. A dedicated user interface is being implemented (see appendix 2 – section 7.3.8) for providing the data subjects with the needed tools for managing consent.

2.2. Consent and the GDPR - brief overview

Consent acquisition will comply with the requirements set out in the General Data Protection Regulation (GDPR). According to GDPR, in order to be valid, consent needs to be (see relevant deliverable D2.1 – section "2.4.1 Consent", for more details):

a) *freely given*: meaning that the data subject cannot be in any way forced or deceived to provide his consent.

b) *Specific*: each consent must be referred to – and so be acquired for – a sole and single purpose (even if this purpose can be relatively broad), allowing the data subject to agree or not with the specific processing operation envisaged by the data controller (the purpose of processing must be specified prior to – and in any event, no later than – the time when the collection of personal data occurs).

c) *Informed*: data subjects will be free to express their consent so long as they are put in the condition to know the needed details relating to the processing of their personal data.

d) *Unambiguous*: this is the main innovation brought forth by the Regulation in connection with the consent. Unambiguous means that the individual's intention to allow the processing of his data must not give rise to any kind of doubt. The relevance of this requirements may be better comprehended in the light of Art. 4, let. 11), of the GDPR, which establishes that the consent can be given «*by a statement or by a clear affirmative action*», hence admitting the so-called "implied consent". In brief, the Regulation allows the controller to construe the data subjects' consent directly from their conduct, to the extent that such actions are unarguably clear and may be demonstrated, if requested.

Art. 7.2 of the GDPR states that if the consent is given in the context of a written declaration which also concerns other matters (e.g. terms of service), the request for consent must be presented in a form that is distinguishable from the rest of the document, in an intelligible and easily accessible form, using clear and plain language.

Regarding consent collected for research purposes, the GDPR acknowledges (recital 33) that "it is often not possible to fully identify the purpose of personal data processing for scientific research purposes at the time of data collection". For this reason, when the consent is asked for these purposes, "data subjects should have the opportunity to give their consent only to certain areas of research or parts of research projects to the extent allowed by the intended purpose".

Finally, since the individual is the pivot of any decision regarding the processing of his personal data, he/she shall have – and be informed about – the right to withdraw his consent at any time, as easily as he originally gave it.

2.3. Summary of the MHMD approach: dynamic consent and smart contracts

To comply with the rules set out in the GDPR, and make it possible to automatically operationalise consent in the context of a blockchain architecture, MHMD will adopt a dynamic consent principle, in consideration of the fact that this principle meets both the needs of hospitals and of citizens, because:

- Dynamic consent is specifically useful when data subjects have to provide directly (i.e., without the intermediation of healthcare professionals) consent for third parties to access their datasets, as it allows citizens to have a clear interface to understand the purpose of data usage and the consequences of the consent, while also selecting privacy and consent preferences in an intuitive and easy-to-understand way.
- Dynamic consent allows hospitals to clearly present relevant information to patients, in order to gather the consent and being sure that the patients have fully understood and agreed to the needed data collection and processing.

At the same time, the dynamic consent will be implemented as smart contracts, thus triggering within the overall blockchain architecture the processes aimed to ensure the traceability of the given consent, the trustworthiness of the data sharing process, and the automatization/operationalisation of the consent preferences, as collected by the hospitals or directly defined by citizens/patients.

Thanks to the smart contracts enacting the "Dynamic consent" principle:

- Hospitals will be able to document the tamper-proof record of the consent obtained from the patient, as to allow an easy traceability of it (also in case of an external auditing procedure), while automating data sharing under specific conditions, providing third parties with ready-to-use consented datasets, without the need for contacting back the data subject, or the data controller itself. "Cleared" datasets will enable easier sharing, thus laying the foundation for a proper health data marketplace.
- Patients will be able to activate directly their data sharing under precise pre-defined conditions. The smart contract will automatically execute the data exchange when the conditions defined in the smart contract will be met by the data access request made by an interested third party (e.g., research centre, pharma industry, etc.).

The following paragraphs will detail the dynamic consent enactment, subsequently explaining its implementation as a smart contract, briefly discussing as well some relevant legal issues.

3. BACKGROUND AND STATE OF THE ART ON CONSENT AND REASONS FOR DYNAMIC CONSENT

The requirement for consent is a fundamental principle of the medical profession and it is tightly linked with legal obligations and ethical principles.

It is a requirement that researchers obtain informed consent from individuals prior to the start of any research. The consent form is the primary means of recording individual decisions about involvement in research, and for determining 'whether additional consent is required, or whether the existing consent covers the new research'.

It can be collected via traditional paper form, or more modern electronic consent forms offered by tools such as REDCap. It forms the contract between patients and researchers, it also outlines how an individual's data will be protected and their privacy preserved.

Research institutions keep samples, such as blood, tissues, organs, etc., donated by data subjects and share them with other research institutes, which use them in medical research. Donors also provide personal information such as date of birth, gender, and sometimes more identifying information such as NHS or other healthcare systems identification numbers. Work is undertaken with a strict ethical code that prevents disclosing any "personally-identifying" data to research institutes. When sharing donors' samples, only a subset of pertinent information about the donors is shared.

Institutions work by collecting samples under a study, and subsequently only share samples with research institutes that are members of this study. Current code of ethics does not allow the use or sharing of collected samples for other than the "original" study and "purpose" they were collected for.

Participants should be informed about the purpose(s) for which their data will (or may) be used; where it will be stored; the expected retention time; if any other parties are involved; the amount and the sensitivity of the information exchanged; whether the data will be shared onward to other parties; whether the consent to use these data can be revoked; the risks and benefits for them of the intended data processing. Consent to data processing is also a requirement of data protection and privacy laws in most countries.

In the case of biobanks, where there are multiple researchers and research projects, obtaining informed consent for all research uses at the time of recruitment can be difficult, and re-consenting for future studies already collected data can be a costly and time-consuming process, with the possibility of stagnating information and incorrect or outdated contact details.

As a solution, the practice of obtaining broad consent has become standard practice in many biobanks, where a person effectively agrees to permit a broad use of their data once it is provided. Consent is still frequently done as a one-off procedure with a paper form for participants to sign; these forms are often lost or filed away, and over time people forget what they have consented to and why. A move to electronic forms has eventually started to spread around, but this still does not negate the latter issue.

In 2014, concerns that the "broad consent" might not remain a lawful option for research, led Jane Kaye and her co-authors to propose a "dynamic consent" approach, whereby, "rather than being restricted to the opportunity only to give broad consent to the use of their samples and data, individuals could provide different types of consent depending upon the kind of study"[3]. These consent preferences would then "travel securely with their samples or data so that third parties know the scope of the consent that applies", and a secure consent interface, tailored to individuals' needs, would "allow participants to change their consent preferences reliably". Dynamic consent – they posited – would meet "the highest international ethical and legal standards for consent in a world where data protection laws are in flux". It would not be meant, therefore, to be "a replacement for existing models such as broad consent, but rather a facilitation tool to improve how that consent is obtained, understood and acted upon". Given an 'opt in' and 'opt out' approach to choose, a participant would still be enabled to "choose to give a broad consent and not receive updates and so on, but if at some future point they [would] wish to become more engaged they [would] have the option to do so"[4].

3.1.THE CASE FOR DYNAMIC CONSENT

MHMD has chosen to explore advanced forms of consent, pursuing in particular the "dynamic consent" approach[5], as a mean to increase the level of awareness and understanding of the data subjects regarding the data usage, while at the same time overcoming the information asymmetries between them and researchers. The concept of dynamic consent followed an intense debate devoted to understanding whether the "'traditional' form of informed consent can and should adapt to emerging forms of research, from genomics and biobanking to increasingly virtual, global research networks assisted by online, openly shared genomic databases"[6], and has been seen as an alternative and a better shaped solution to overcome the possible weaknesses of broad consent, whenever its very nature of informed consent or real consent happens to prompt critical appraisals.

The approach of integrating a Dynamic Consent system to existing systems implies a willingness to modify the relationship between donors, institutions, and researchers - from a passive relationship where the donor role ends by signing long consent forms and donating samples to a more interactive relationship by bringing the patient in the loop. Applying dynamic consent gives the donor the opportunity to view how their samples are used and shared, by whom, and the impact of their contribution to research. It also allows the donor to widen the use of their samples to other research fields, purposes, or other studies. At the same time, the data subject can create black lists to exclude certain research institutes from using their samples, or to

[3] J. Kaye, E.A. Whitley, D. Lund, M. Morrison, H. Teare and K. Melham, Dynamic consent: a patient interface for twenty-first century research networks, "European Journal of Human Genetics", 2014.

[4] E. Morley-Fletcher, Enhanced Consent: a vision for Patient Data Protection and Data Management, Paper for Networking session held within the ICT2015 event in Lisbon, Portugal.

[5] Kaye, J., Whitley, E. A., Lund, D., Morrison, M., Teare, H., & Melham, K. (2015). Dynamic consent: a patient interface for twenty-first century research networks. European Journal of Human Genetics, 23(2), 141-146.

[6] B. Schmietow, Ethical Dimensions of Dynamic Consent in Data-Intense Biomedical Research—Paradigm Shift, or Red Herring? in D. Strech, M. Mertz, (Eds.), Ethics and Governance of Biomedical Research - Theory and Practice, Springer International Publishing Switzerland 2016.

exclude contributing to certain research fields. It is because the donor is given the chance to change and modify his consent on-line, tracing their samples and data, as well as viewing the effect of their consent, that the term dynamic consent is given.

This approach follows the more general pursuit of patients and citizens empowerment, as advocated by several European initiatives, toward the implementation of patient-centred and participatory medicine, from the foundational eHealth Action Plan 2012-2020[7], to the most recent "Digital Health Society Declaration"[8], passing through independent initiatives such as the MyData movement[9].

Following these initiatives, in particular when individual citizens' direct data sharing is concerned, one of the objectives pursued by MHMD is in fact to empower citizens with dynamic consent mechanisms and gain their trust with a transparent solution that protects and manages their personal information. In this respect, Dynamic Consent processes within MHMD aim to assure that:

- Data subjects are offered information with regard to how their information is used, for which specific purposes, and by whom.
- Data subjects are offered controls to change their mind by changing their consent when they learn about how their information is used.
- Consent is enforced within the law and within specific organisational policies.

3.2. Key features of Dynamic Consent

Dynamic consent can be described as "an interactive interface that allows participants in research to choose and alter consent choices in real time. The system provides reliable storage and enforcement of these choices by cryptographically protecting sensitive personal information in a way that allows data to be accessed in only those ways for which consent has been given"[10], also allowing for tailoring consent "on a wider variety of research initiatives, in a more open and more flexible manner"[11].

Dynamic consent aims to bridge the gap between informed consent and broad consent, providing a digital interface from which patients can engage researchers, define their consent profile, and modify their consent profile over time. Patients can then start from the premise of the broad consent model and visualise

[7] E-health Action Plan 2012-2020 - Innovative Healthcare for the 21st Century - Communication from The Commission to The European Parliament, The Council, The European Economic And Social Committee And The Committee Of The Regions, Brussels, 6.12.2012

[8] https://www.eu2017.ee/news/insights/digital-health-society-declaration

[9] https://mydata.org/declaration/

[10] J. P. Woolley, How Data Are Transforming the Landscape of Biomedical Ethics: The Need for ELSI Metadata on Consent, in B. Mittelstadt, L. Floridi (Eds.), The Ethics of Biomedical Big Data, Springer International Publishing Switzerland 2016.

[11] B. Schmietow, Cit.

how their data is being used by researchers, and hence tailor their future consent parameters to conform to their own views and wishes.

This is what makes the consent ‘dynamic’, as it allows interactions over time, and enables participants to consent to new projects or to alter their consent choices in real time as their circumstances change, and have the confidence that their changes take effect, as they are able to see how and where their data is used, and for which studies.

Rather than being restricted to give broad consent at the moment their data and samples are provided, patients can define different types of consent depending upon the kind of study, the purpose of the study, or the organisation leading the study.

An electronic interface would also allow the patient to alter their contact details, so that data obsolescence becomes less of a problem, and allow the provision of additional data from other sources, such as the use of wearable technology, allowing the patient to feed additional data back should they desire.

Such an interface then allows patients to engage in their own time, as much or as little as they choose. It could also allow them to receive information on the use of their samples and data, enrol in new studies, and complete further surveys, all of which could drive their consent choices, and make them far more engaged in the research activity.

The dynamic consent approach provides the following key advantages:

- It meets the highest international ethical and legal standards for consent in a world where data protection laws are in flux.
- It enables participants to keep all of their information in one place, with a record of consent and research involvement, thus enabling more active engagement in research.
- Collection of one-off consent for research can often occur at a stressful time for the person concerned, such as before treatment or surgery; dynamic consent removes this pressure by allowing participants to return to their decisions and review their consent preferences in their own time.
- It promotes scientific literacy as participants become more informed about the research carried out on their samples and information, which encourages public trust by making research more transparent and accountable.
- For researchers, it provides an easy mechanism to identify individuals who have consented to being approached and recruited for new studies, to participate in online surveys or to canvas opinions.
- It can be tailored for specific situations, as a ‘one-stop' interface to facilitate better translational research and to coordinate clinical and research activities co-ordinated around the patient

3.3. ELEMENTS OF A DYNAMIC CONSENT FORM

As useful reference for the identification of the operational clauses to be included in the dynamic consent form and subsequent smart contract, the `data licensing agreement' (DLA)[12] is being considered, in view of exacly reflecting the prescriptions of Art. 13 of the GDPR. Such DLA aims at making sure "that data are only processed insofar as necessary for the performance of the agreement by a party that is not allowed to share the data with other parties" (thus, only covering one of the possible options contemplated within MHMD). The DLA provides an interesting example of how to put in place a proper user interface, making clear the fundamental aspects of the data sharing process to the data subject, following an approach already developed within the dynamic consent implementation: each clause is displayed in the interface on a separate screen, thus making the DLA easily readable, also making use of animation to simplify the content and highlight the key aspects to be taken into account by the data subject.

The DLA is composed of general clauses and modular clauses, that are presented as individual items below. These clauses are useful examples of operational clauses to be included in the dynamic consent management interface and translated in actionable code for the smart contract.

General clauses

1) identifying the parties to the contract:

(1) the data subject: a patient or e.g. a user of a health App;

(2) the data controller(s): an identified health-App service provider, doctor, medical specialist or e.g. a hospital, insurance company, research institute or pharmaceutical company.

The data subject:

licenses the identified data controller(s) who is (are) a party to the DLA to use (process):

{ a specified set (stream) of her or his personal data;

{ for explicitly specified purpose(s);

{ clearly expressing unambiguous and informed consent for the processing of his or her sensitive data for the explicitly specified purpose(s).

The data controller(s) will use (process) the data:

{ only for the specified purpose(s) and | if necessary | for purposes that are deemed compatible (no re-use out of context);

[12] Verheul, E. R., Jacobs, B., Meijer, C., Hildebrandt, M., & de Ruiter, J. (2016). Polymorphic Encryption and Pseudonymisation for Personalised Healthcare. *IACR Cryptology ePrint Archive*, *2016*, 411.

{ whenever permitted, relying on a valid legal basis alternative to the data subject's consent (e.g. when the data have been manifestly made public by the data subject), by making very clear, in case, the existence of a legitimate interest which is not overridden by the interests or fundamental rights and freedoms of the data subject [as per Art. 13.1, let. (c) and (d) of the GDPR];

{ employing additional techniques of anonymisation and pseudonymisation (if the data enables re-identification because it is linked with other data, or if it is unique within the dataset);

{ deleting the data once the purpose is exhausted;

{ but always within a specified time period (which can be extended with a renewal of the DLA if the purpose has not yet been exhausted);

{ confirming that the data subject has the right to withdraw her consent at any time (which only regards future processing);

{ providing an easy way to withdraw consent;

{ implementing measures as are needed to ensure that the data subjects may exercise their right to request access to and rectification, erasure and portability of personal data, or to object or ask restriction of processing concerning them [as per Art. 13.2, let. (b) of the GDPR] ;

{ providing an easy way to receive an electronic copy of the data processed;

Modular clauses

the modular clauses might include (but are not limited to) the following aspects:

- specify the identity of the data processor(s), and/or
- specify the recipients or categories of recipients of the personal data [as per Art. 13.1, let. (e) of the GDPR]; and/or
- specify whether data may be processed outside the EU (based on what legal safeguards), and/or
- stipulate with whom the abstract results (which are not personal data) may or may not be shared, notably whether or not these abstract results may be shared with commercial companies, and/or
- specify the type of analytics that will be employed [as per Art. 13.2, let (f) of the GDPR: the data subject must be informed about "*the existence of automated decision-making, including profiling, referred to in Article 22(1) and (4) and, at least in those cases, meaningful information" must be provided by the controller "about the logic involved, as well as the significance and the envisaged consequences of such processing for the data subject*"].

4. DYNAMIC CONSENT AS A SMART CONTRACT

As briefly introduced above, MHMD changed its initial approach, which involved the simple submission of the Dynamic consent to the local Ethics Committees, to implement it as a proper smart contract, to be fully integrated with the overall MHMD blockchain architecture.

The key concept laying behind the implementation of the dynamic consent as smart contract is mainly to ensure the trustfulness of the data sharing, identifying the data subject and his/her consent preferences, without needing to rely exclusively on the assumed trusted third-party/data controller preliminary acquisition of such consent, this way providing to the data controller a transparent and tamper-proof audit trail mechanism, which makes it possible to verify whether a proper consent is in place, and if the relevant regulations are effectively respected. At the same time, this approach provides the individual user, when directly sharing data, with an easy to understand interface for selecting the consent features (purpose, time of availability of the dataset, subjects allowed to access), which are then “translated” into a smart contract that executes automatically data transactions when the consent conditions are met in the data access request.

The “translation” process is not an easy task, as it requires that the consent preferences are transformed in code able to trustfully reflect those preferences, enabling the desired actions without putting at risk the privacy of the data subject or the security of the dataset, and not allowing non-consented activities over the data and/or non-consented access. The following sections will provide an overview of the current challenges and approaches, also indicating MHMD choices and subsequent activities to address above-mentioned issues and effectively manage the dynamic consent through smart contract.

4.1. OPERATIONALISE CONSENT – AN OVERVIEW

Once the data subject has been allowed to define, thanks to dynamic consent interfaces on the mobile application (see the preview on p.11, the full interface design will be provided in D3.3), at a high level of granularity the terms under which his/her personal sensitive datasets can be accessed and used, the next challenge is to make this consent, privacy preservation solutions, and relevant exchange protocols, available as machine readable expression in order to enable a secure and trusted data sharing environment[13]. The same applies in the case of data controllers willing to have their available datasets pseudonymised and shared with third parties for research purposes and other consented usages, or anonymised.

[13] An emerging offer of "Personal information management services" - Current state of service offers and challenges - European Commission Position Paper, January 2016.

This is particularly useful whenever a data platform is willing to enable automatically micro-payments and/or service provision to the individual or institution in exchange for data[14], thus making personal health data a valuable digital asset[15].

At the same time, there is a need for making the data exchange lawful, taking into account the relevant regulations, in an automatic fashion, thus without requiring (or at least minimising) the need for assistance or ex-post enforcement control by a trusted third party. Recent techniques, such as Data provenance (i.e., watermarking privacy preferences into the data), allow to attach information about the provenance of the data to the data itself in a way that cannot be tampered with[16].

MHMD has explored various initiatives addressing the above-described challenges in consent operationalisation, as necessary preliminary steps toward the design and implementation of a proper solution for operationalising the dynamic consent approach, aligning this process with the relevant regulation, and at the same time minimising the risk of unintended activities over the shared dataset. The following are the current available solutions taken into account:

- Meeco[17], has implemented a Permission Consent Management layer, which enables users to indicate their preferences regarding terms and duration of data sharing, while also making it possible to update the terms, deleting the data, or revoking the consent.
- The OpenConsent group[18] proposed the concept of "consent-as-a-service", enabling consent tracking and also the development of standard consent receipts[19].
- Stemming from the cooperation among COALA IP[20] (Coalition of automated legal application) and Global Consent[21], the Smart Consent Protocol[22] explored the extension of the Consent Receipt approach, developed within the Kantara Initiative[23] (also partner of the OpenConsent Group). The Consent receipt "records a standard set of legal, social, and contextual parameters relating to an

[14] Ibid.

[15] Following also an approach suggested in S. Conway, L. Spies, J. Endersby, T. Daubenschütz, *Smart consent protocol - A White Paper from Rebooting the Web of Trust III Design Workshop*, March 2017. https://github.com/WebOfTrustInfo/rebooting-the-web-of-trust-fall2016/raw/master/final-documents/smart-consent-protocol.pdf

[16] Ibid.

[17] https://meeco.me/

[18] https://openconsent.com/

[19] An emerging offer of "Personal information management services"…, op.cit.

[20] http://coala.global/

[21] http://www.consent.global/

[22] *Smart consent protocol - A White Paper from Rebooting the Web of Trust III Design Workshop*, op.cit.

[23] https://kantarainitiative.org/

information-sharing transaction"[24], enabling people to clearly define their privacy and data sharing preferences, and thus increasing trust in the overall system. The relevant Consent Receipt Specification[25] are currently in preparation, building on top of the ISO 29100 Privacy Framework common privacy terminology, and designed to be compliant with the most recent ISO 29184 - Guidelines for online privacy notices and consent, currently under development.

- The above-mentioned Smart Consent approach proposes the integration of the relevant data Terms of Use within the Consent Receipt, and then the storage of such receipt "with an immutable proof (hash value) that cannot be repudiated, on a decentralised public ledger"[26]. This can be obtained by implementing the Smart Consent as a Smart Contract.
- Again, within the Kantara Initiative, the User Managed Access (UMA) has been implemented, consisting of a "OAuth-based protocol designed to give an individual a unified control point for authorizing who and what can get access to their digital data, content, and services, no matter where all those things live"[27], enabling federated and consent-driven data transactions.
- Building on top of the UMA protocol, HIE of One (Health Information Exchange of One)[28] is exploring the self-sovereign technology[29], defined as a technology that is not under the control of any institution[30], and that acts as one person's agent or fiduciary, designed to issue authorization tokens for access to personal data directly from compatible resource servers in other domains. The system works upon a dedicated authorization server (AS), based on the UMA standard, which issues authorization tokens based on the subject's policies and the claims submitted by would-be data users. The AS benefits also from Artificial Intelligence, through which it can learn and updated the data subject policies.

As a result of the above-reported overview of the current state of the art, MHMD is working on the combination of various of the above-mentioned features in order to:

- Provide the patient with an easy-to-use interface for managing his/her data, and expressing his/her consent options, monitoring the activities on the shared data, and update the terms of consent over time or revoke it. The mobile application prototype, which will be developed by HWC

[24] *Ibid.*

[25] https://kantarainitiative.org/confluence/display/infosharing/Consent+Receipt+Specification

[26] *Ibid.*

[27] https://kantarainitiative.org/confluence/display/uma/Home

[28] http://hieofone.org/ - White paper at: https://drive.google.com/file/d/0B_Rve6d1gHMWNmtueU5mTm9wVEk/view

[29] P. Sheldrake, Defining Sovereign Technology, so we can build it, and so we know it when we see it – Medium, May 26, 2016 https://medium.com/@sheldrake/defining-sovereign-technology-so-we-can-build-it-and-so-we-know-it-when-we-see-it-98ad77914025

[30] HIE of one White paper, op.cit.

(D3.3, due at M16) will encapsulate these principles, designing an easy-to-use interface for defining consent options and other data sharing policies.

- Embed into the relevant smart contract both operational and non-operational clauses (see the dedicated section below), similarly to the Consent Receipt approach, in order to trigger the data processing only on specific user-defined conditions, guaranteeing at the same time its lawfulness, and including in the consent smart contract both specific (actionable) preferences and options and relevant regulatory and legal framework of reference. This activity will be performed after the first release of the smart contract template and relevant preliminary code (due by M24), as the relevant legal prose and parameters have to be preliminarily defined (also on the basis of the conclusions which will be provided in D2.2 Legal Opinions on the Project Assessment– due in M15 - and D2.3 Study on Dynamic Consent – due in M16).
- Using a learning algorithm to allow patients/citizens to set at the beginning a "tree of decisions/preferences" to be automatically/semi-automatically translated into the consent options applicable to specific real-world scenarios, reducing the need for direct intervention of the user himself. This feature will be further explored and embedded in the final version of MHMD smart contract's implementation, due by M36.

4.2.SMART CONTRACTS – GENERAL OVERVIEW

Originally popularised in 1994 by Nick Szabo, the smart contract has been defined as "a computerized transaction protocol that executes terms of a contract". The technology, though, was not fully developed then, because of the relevant technical issues. In fact, although smart contracts might have been technically feasible and thus adoptable in several use cases since the mainstream availability of personal computers and internet, some practical issues have been raised, from the need of programming and running the code separately, to the need of trusting/relying on the other party's code, as well as the risk of having different outcomes from the implementation of two different instances[31].

Distributed ledger technologies – brief definition

A Distributed Ledger Technology (DLT) can be defined as "computer software that is distributed, runs on peer-to-peer networks, and offers a transparent, verifiable, permanent transaction management system maintained through a consensus mechanism rather than by a trusted third-party intermediary, and that guarantees execution" [as defined in Reyes, Carla L., Conceptualizing Cryptolaw (February 9, 2017). Nebraska Law Review, Forthcoming. Available at SSRN: https://ssrn.com/abstract=2914103].

A DLT is distributed because "the record is held by each of the users (or nodes) on the network and each copy is updated with new information simultaneously", thus being able to maintain – without need of reconciliation – only one record which represents "a golden source of data" [ISDA, op. cit.].

For this reason, Smart Contracts have been revived by the advent of Blockchain, and distributed ledger technologies (DTLs – see box) in general, whereas the smart contract – once embedded in the distributed ledger – becomes the only valid version of the agreement between the parties, impossible to be tampered by any of them, and also executes itself automatically (is self-enforcing), without requiring any intervention by any of the stakeholders, nor any form of mutual trust in the correct execution of the agreed-upon contract[32].

In the context of DLT, a smart contract is "a computer protocol—an algorithm—that can self-execute, self-enforce, self-verify and self-constrain the performance of a contract"[33]. As integration of this definition, it can also be added that "[a] smart-contract is an event-driven program, with state, which runs on a replicated, shared ledger and which can take custody over assets on that ledger"[34].

[31] As explained in ISDA Whitepaper Smart Contracts and Distributed Ledger – A Legal Perspective, ISDA Linklaters, August 2017.

[32] Ibid.

[33] Swanson, T. (2014). Great Chain of Numbers. *A Guide to Smart Contracts, Smart Property, and Trustless Asset Management, publisher= Creative Commons-Attribution 4.0 International*.

[34] R. G. Brown, *A simple model for smart contracts,* February 2015, http://gendal.me/2015/02/10/a-simple-model-forsmart-contracts

The following are the key characteristics of a smart contract[35]:

- Smart contracts are software programs that run on DLT;
- Smart contracts offer event-driven functionality—when triggered by external data (which might be provided by "oracles", i.e. trusted data sources that send information to smart contracts), smart contracts will modify other data;
- Smart contracts can, acting on information provided by oracles, enforce a functional implementation of a particular requirement, and can show proof that certain conditions were met or not met;
- Smart contracts can track changes in state over time;
- Smart contracts are autonomous in that the software developer who created them need not to actively maintain, monitor, or even be in contact with them while they operate;
- Smart contracts are distributed because they exist as software running on a DLT protocol that itself is distributed across a variety of network nodes;
- Smart contracts guarantee execution of the contemplated transaction.

Although these elements provide a general definition of a smart contract, the current debate transcends this technological characterisation, and focuses more on real-world application scenarios for this innovative technology. In such a context, the term smart contract might indicate two different things:

1) the first one identifies a specific technology: "code that is stored, verified and executed on a blockchain"[36], and which is usually referred as "smart contract code". This smart contract code indicates proper software agents "fulfilling certain obligations and exercising certain rights, and may take control of certain assets within a shared ledger"[37]. Being executed by the blockchain, this contract "will always execute as written and no one can interfere with its operation"[38]. Smart code contracts are typically involved in addressing only operational aspects of a contract/agreement.

2) the second one refers to a specific use case for smart contract code, to provide a "way of using blockchain technology to complement, or replace, existing legal contracts"[39], thus addressing the issue of "how legal contracts can be expressed and implemented in software"[40]. This sort of contracts will therefore combine code and traditional legal language, encompassing both operational and non-

[35] As summarised in J.D. Hansen, Carla L. Reyes, *Legal Aspects of Smart Contract Applications – Perkins & Coils White paper*, Perkins Coie LLP, May 2017.

[36] J. Stark, *Making Sense of Blockchain Smart Contracts*, CoinDesk, June 2016 https://www.coindesk.com/making-sense-smart-contracts/.

[37] Clack, C. D., Bakshi, V. A., & Braine, L. (2016). Smart contract templates: foundations, design landscape and research directions. *arXiv preprint arXiv:1608.00771*.

[38] J. Stark, op. cit.

[39] Ibid.

[40] Smart contract templates, op. cit.

> operational aspects (such as how to interpret the legal prose), "some of whose operational aspects must be automated[41] (using smart contract code).

Coming to the distinctions between non-operational and operational aspects of a contract, the following explanations are needed:

- Operational aspects are "the parts of the contract that we wish to automate [involving] precise actions to be taken by the parties and therefore are concerned with performing the contract]"[42].
    - These clauses usually embed "some form of conditional logic" which at the occurrence of a specified event trigger a deterministic action.

- Non-operational aspects refer to those parts of the contract that the parties "don't wish to (or cannot) automate" [43], and relate to the "wider legal relationship between the parties"[44].
    - Non-operational clauses might include[45]:
        - A clause specifying what law should govern in the event of any dispute;
        - A clause specifying what jurisdiction any disputes may be brought in;
        - A clause providing that the written legal document represents the entire agreement between the parties;
        - A representation that a party's obligations under the legal agreement constitute legal, valid and binding obligations.

Following the two definitions above, and in order to achieve an abstraction level capable of encompassing both of them, the following definition has been proposed:

> "A smart contract is an automatable and enforceable agreement. Automatable by computer, although some parts may require human input and control. Enforceable either by legal enforcement of rights and obligations or via tamper-proof execution of computer code"[46].

[41] Ibid.

[42] Ibid.

[43] Ibid.

[44] ISDA Smart Contract White Paper

[45] Ibid.

[46] Ibid.

4.3. MyHealthMyData Consent smart contract

The following scheme provides a simplified overview of the MHMD architecture and the relevant data flow (for a more detailed and accurate description, please refer to D6.3 and subsequent updates).

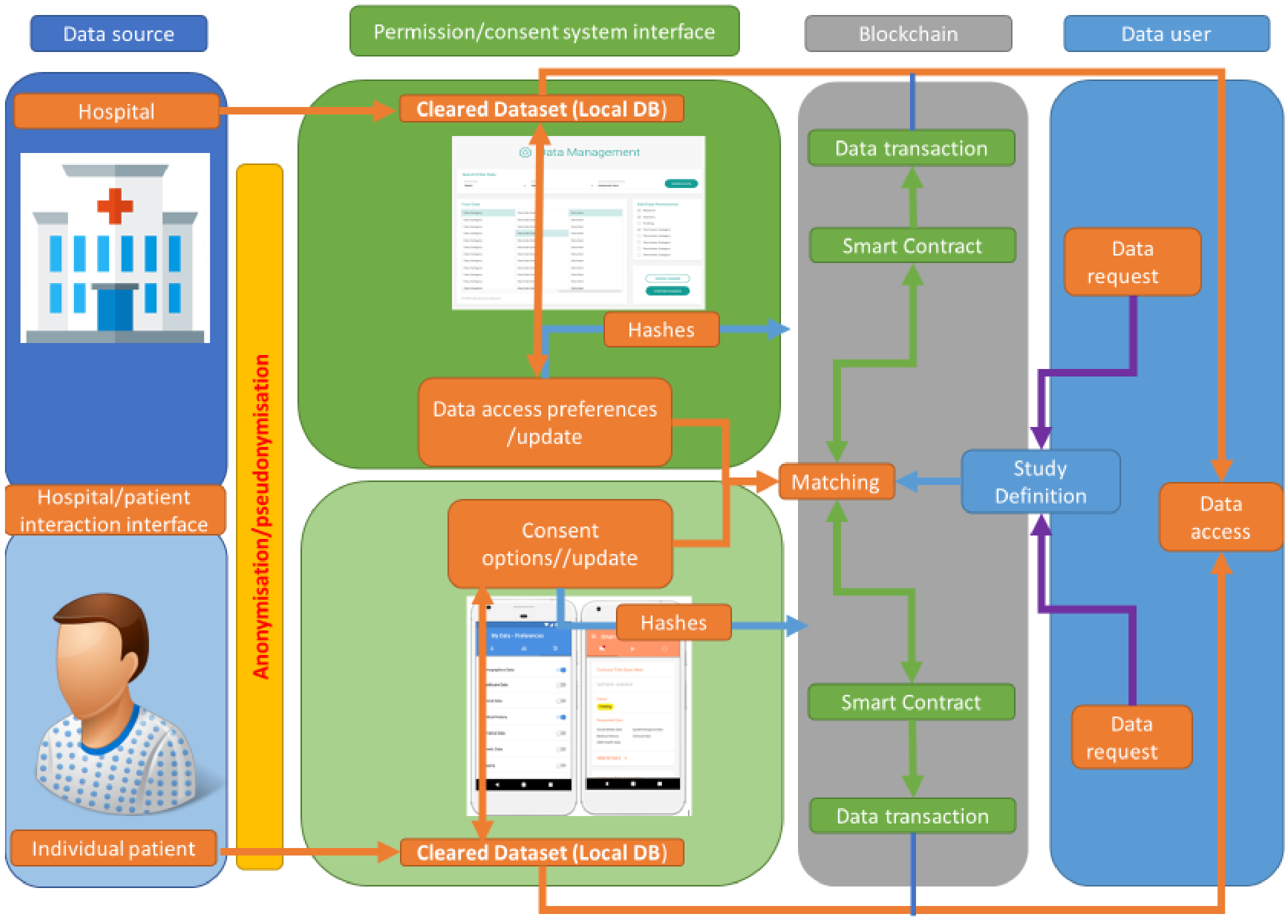

Hospitals and patients can share cleared datasets (previously anonymised/pseudonymised), associating to those datasets the relevant permission/consent options. Through the dedicated interface (web application for hospitals and mobile application for patients), the data subject can specify the chosen data access policy with a high level of granularity, defining allowed purposes of the data usage (e.g., only for public research, only for specific disease, etc.), time of availability (data available until…), specific data users (e.g. only non-profit institutions), type of analysis permitted, etc. Through the same interface, these access policy/consent preferences may be updated in time. Hashes of the consent – generated automatically – are uploaded on the blockchain, to ensure that the consent options are not tampered afterwards.

On the data user side, a data study request is shared on the blockchain, defining data type, purpose of the study, and other usage condition.

The API (which incorporates the Permission/consent system) automatically verifies the compatibility of the data request with the defined consent options, without sharing any information about the data.

At this point, the smart contract (which "lives" in the blockchain), and which is already informed of both the existence of the consent and of the data request, is informed that all conditions for allowing the data exchange are matched, and thus triggers a new transaction enabling the actual data exchange, at the same time leaving track of the transaction on the blockchain itself (thus ensuring the lawfulness of transaction, and its auditability in the future). Only at that point, after having obtained the authorisation from the relevant smart contract, access to the relevant cleared datasets (through the cleared local database, not to be confused with the internal dataset of the hospital – which is never accessible from outside) is granted to the data user for performing its study. For obvious reasons (i.e. the cleared datasets are not stored on the blockchain) the actual data transfer takes place offchain, in MHMD cloud environment.

If one of the conditions needed for allowing the data exchange are not yet met, the smart contract keeps itself on hold, waiting for all conditions to be verified. A number of pre-defined smart-contracts might be shared on the blockchain for specific kinds of data (e.g., anonymous, pseudonymous, for consent to specific research purposes, for broad consent to research activities in a given field, etc.).

MHMD initially agreed on a preliminary template for a data-exchange smart contract, which would encompass the (dynamically managed) consent options, identifying some key variables that will be included:

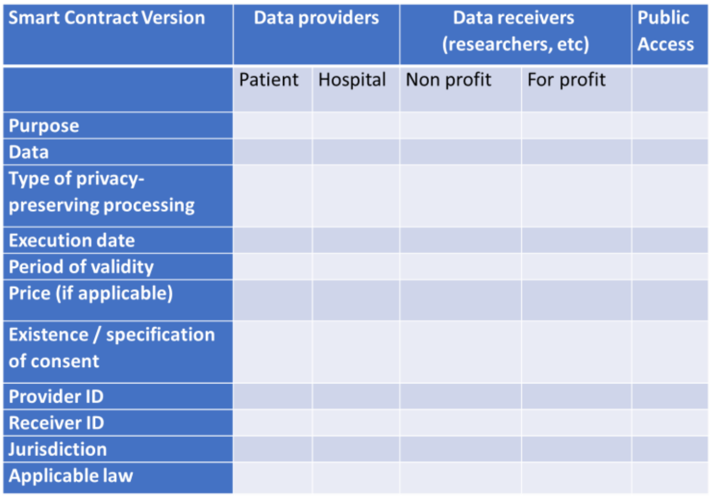

Smart Contract Version	Data providers		Data receivers (researchers, etc)		Public Access
	Patient	Hospital	Non profit	For profit	
Purpose					
Data					
Type of privacy-preserving processing					
Execution date					
Period of validity					
Price (if applicable)					
Existence / specification of consent					
Provider ID					
Receiver ID					
Jurisdiction					
Applicable law					

Eventually, however, the smart contract template will need to reflect, in order to properly match the individual's data processing and sharing permissions (based on very granular choices) and stakeholder's access requests, almost all the elements included in the DLA quoted above.

The very essence of the agreement in place between the data subject and the data controller is the privacy information notice plus the consent provided accordingly, whenever required, and/or the other valid legal basis which may substitute consent (possible serving as an equivalent encoded trigger). Only the information notice includes all the elements which are needed, under the applicable data protection law, in order to verify whether a certain data processing operation has been carried out or not in compliance with the requirements in force.

The nature of the variables is broad: from simple values (e.g., for data and IDs), to strings (e.g., for applicable law and jurisdiction), to more semantically complex operations defining the consent options and preferences. Also, some of the variables are computable operational clauses (as per the definition above), like price, execution date, period of validity: these clauses pertain the execution of the contract and need to be automated. Other variables, such as purpose, applicable law, jurisdiction, are clearly non-operational, as they present rather large and highly-complex semantic structures to express/explain relevant rights and obligations in case of non-compliance, or to provide information about governing law, jurisdiction, etc.

Thus, the issues of how to implement all these different operational and non-operational clauses, and how to "translate" and "automate" most of these clauses within a smart contract, will be properly addressed in MHMD.

The following paragraph presents the currently available solutions for this issue.

4.4.SMART CONTRACTS: HOW TO REPRESENT THE SEMANTIC RICHNESS

4.4.1. RICARDIAN CONTRACT

According to the relevant foundational paper, a Ricardian Contract can be defined as:

"a single document that is a) a contract offered by an issuer to holders, b) for a valuable right held by holders, and managed by the issuer, c) easily readable by people (like a contract on paper), d) readable by programs (parsable like a database), e) digitally signed, f) carries the keys and server information, and g) allied with a unique and secure identifier"[47].

This means that the same document can be read by lawyers or contracting parties, which would be able to understand the key elements of the contract, while at the same time can be machine-readable, i.e. expressed and executed in software[48].

This means that the legal provisions set forth by the privacy notice (constituting the perimeter of any processing of data), together with the specific processing-triggers arising from the consent(s) given by the data subject (or other applicable legal basis) can be embedded in a Ricardian contract, which

[47] I. Grigg. The Ricardian Contract. In Proceedings of the First IEEE International Workshop on Electronic Contracting, pages 25-31. IEEE, 2004. http://iang.org/papers/ricardian_contract.html.

[48] Smart Contract Templates, op. cit.

reflects, in a mostly legal prose document, a mark-up language able to represent the key parameters of the agreement, which are appropriately automated[49].

The intersection between Ricardian and Smart Contract has been explored: whereas smart contracts (in the original Szabo definition – see above) are mainly meant to automate a performance when the contract has been already agreed by the parties (i.e., a "design to capture the flow of actions and events (e.g., delivery of payments) within the performance of a contract"[50]), the Ricardian contract, being "conceptually unlimited in the richness of semantics"[51], is the most appropriate tool for capturing the intent of the parties within the agreement, before its execution (i.e., "captures the meaning of the flows in a way that is secured to your actions within the contract"[52]).

For this reason, it has been said that "the smart contract and the Ricardian Contract are therefore doing different parts of the same process"[53], whereas the latter provides semantic richness representation capabilities, while the former delivers operational performance (as per the following graph[54]):

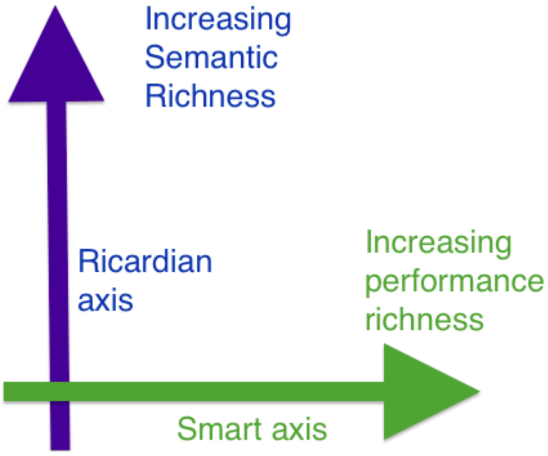

For this reason, an appropriate integration between the two tools has been advocated, in order to add the semantic richness of legal documentation to smart contract services[55].

[49] https://en.wikipedia.org/wiki/Ricardian_contract

[50] Grigg, I. On the intersection of Ricardian and Smart Contracts, http://iang.org/papers/intersection_ricardian_smart.html

[51] Ibid.

[52] Ibid.

[53] Ibid.

[54] Ibid. Legal semantics *versus* operational performance.

[55] Ibid.

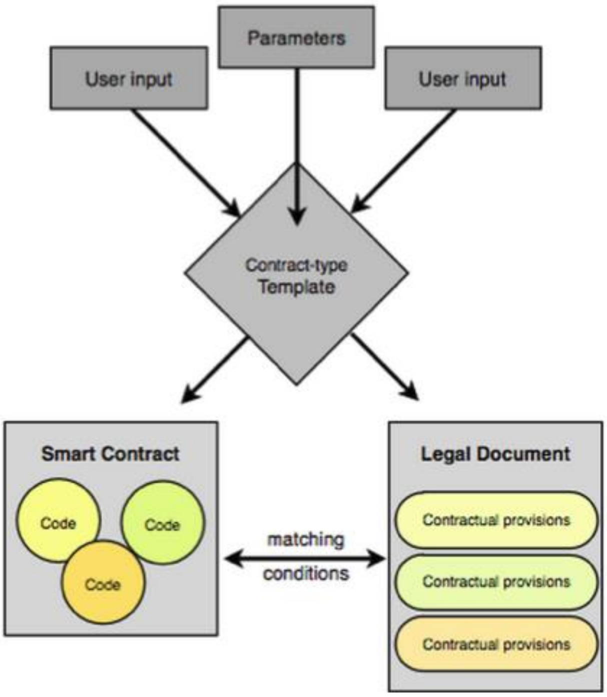

This integration between the two tools will involve the extension of the Ricardian Contract form to refer directly to code by means of hashes: this explicit referral can "pass legitimacy from over-arching legal prose to the code"[56], giving birth to a hybrid form of Ricardian contract based on the triple of "prose, parameters and code": "The legal prose is linked via parameters (name-value pairs) to the smart contract code that provides automation"[57].

4.4.2. SMART CONTRACT TEMPLATE

Building on the concept of Ricardian contract, the Smart contract template approach has been developed[58], with the idea of providing a "framework to support complex legal agreements for financial instruments, based on standardized templates"[59]. This happens using parameters to connect "legal prose to the corresponding computer code, with the aim of providing a legally-enforceable foundation for smart legal contracts"[60].

[56] https://en.wikipedia.org/wiki/Ricardian_contract

[57] Smart Contract Templates, op. cit.

[58] Ibid.

[59] Ibid.

[60] Ibid.

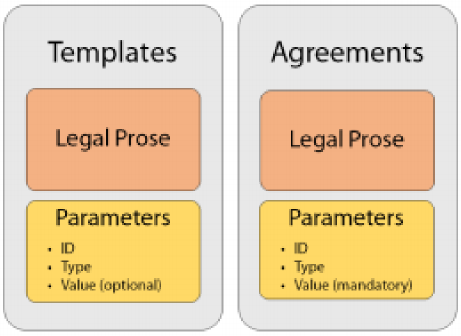

Operational parameters can therefore become an extension of the legal documentation, aimed at directing smart contract code's actions. The templates thus allow the creation of "legally-enforceable smart contracts as electronic representations of legal documents containing prose and parameters. Agreements are then fully-instantiated templates, including customised legal prose and parameters"[61].

Three future challenges will play a key role in ensuring the success of this approach:

- To increase parameter sophistication, to include higher-order parameters, to encapsulate specific concepts usually expressed in legal prose, thus integrating formal logic in legal prose and make it admissible in courts.
- To increase the use of common standardised code, to increase adoption by adopting e.g., common utility functions, able to encapsulate generic parameters/agreements and are identical among all counterparties.
- To create (in the far future) a new formal language (source language), which can be translated automatically both in legal prose and code, directly admissible in court.

4.4.3. MONAX'S DUAL INTEGRATION

Starting from the consideration that – at least in the near future – it is unlikely that courts (and the legal systems at large) are going to accept code as the only mean/source for resolving disputes stemming from smart contract execution, the relevant prose agreement is still central for solving emerging disputes. In order to avoid mismatches between the intention of the parties who have agreed on a given smart contract and the legal prose used by Courts to sort out the outcome of a dispute, Monax has proposed another solution aimed at bridging the gap between existing electronic contracts law and blockchain smart contracts: *dual integration*[62].

[61] Ibid.

[62] https://monax.io/explainers/dual_integration/.

This consists in the integration of a specific legal contract into a corresponding smart contract running on a distributed data store, such as Monax. This solution allows the agreeing parties to use "established dispute resolution processes in the jurisdiction(s) of choice while also using a smart contract as the primary mechanism for administering the data-driven interaction that attends to the agreement between the parties"[63].

4.4.4. LEGALESE

Special mention has to be made for Legalese, which is "an open-source computational legal project to draft legal documents the way programmers develop software"[64]. The basic idea is that of implementing a "domain-specific language (DSL) for legal that is specifically designed to capture legal semantics and logic". Their current R&D activities are focusing, among other things, on smart contracts.

4.4.5. MHMD APPROACH

As a result of the above-described overview, MHMD will follow a threefold approach:

- On one hand, MHMD will identify – starting from the dynamic consent form - the set of operational clauses that will have to be available within the user UI (both in the hospital and in the individual citizen/patient case) and be represented in the smart contract as a set of conditional actions, executed when the data access request matches the data terms of use.
- On the other hand, also on the basis of the already completed review of the relevant regulatory and legal framework, MHMD will identify the high-level legal prose, which has to be embedded within the smart contract in order to ensure the lawfulness of the transaction and guarantee all relevant stakeholders. A particular effort will be devoted to the GDPR and to its key articles regarding consent, data sharing, and research activities. The final aim will be to have a set of parameters representing high-level GDPR's provisions (legal prose), ready to be embedded in the smart contract.

At the end of this procedure, the most suitable technical methodology for representing such legal prose in the smart contract will be selected, taking in consideration the smart contract template approach and the DLA. As a general principle, MHMD will seek to re-use as much as possible standardised codes and procedures,

[63] Ibid.

[64] From the website: http://legalese.com/.

while making publicly available, at the same time, the implemented solution, also in view of enabling an external auditing of the developed solution. As per the project plan, a first smart contract implementation will be ready within M24 (October 2018). To meet the deadline, MHMD will take care both of the development of the dynamic consent form/interface and of the parameterisation of the relevant legal prose, with the aim of having the list of operational clauses and high-level legal parameters ready within June 2018).

The present document will be updated to include the first complete design of the user interface for managing consent (which will be provided by M16, through D3.3), and the underlying technical architecture, as well as a list of operational clauses and associated actions, for external auditing.

The subsequent update (expected to be ready by the end of Y2 – October 2017) will also include the translation into smart contracts, providing evidence of the link between user-selected clauses and subsequent actions expressed in code, which will be triggered within the smart contract when the specific conditions required by the data terms of use will be matched.

5. KEY LEGAL ISSUES

5.1. CONSENT GENERAL ISSUES

A number of issues have been identified by MHMD as part of the preliminary analysis of the regulatory framework, and within various technical workshops organised during the first year of activity.

Particular attention has been devoted to the concepts of anonymisation and pseudonymisation, processing purpose, secondary use, encryption, and their connection with the data subject's consent and the hospital permission system. The issues presented below are taken into account by MHMD for the smart contracts technical implementation, whereas, depending on the "legal status" of the data object of a given transaction, the smart contract might automatically trigger specific actions, allowing (or denying) access to data, or specific data processing activities.

The sections below address separately the identified issues, providing a legal opinion on each of them, with the aim of informing the subsequent MHMD approach in designing the smart contract logic/architecture.

Anonymisation and pseudonymisation: is there a need for consent?

Anonymisation and pseudonymisation do not require the data subjects' consent, as they do not amount to an autonomous purpose for collecting and processing the data. Rather, de-identification techniques (namely, anonymisation and pseudonymisation) constitute a mean to rely on, an intermediate step to be implemented, in order to achieve the real envisaged purpose, such as sharing the anonymized data with third parties for extra-research objectives in absence of the data subjects' consent. For this reason, the application of de-identification measures, as from-time-to time evolving based on daily technical developments, does not need to be grounded on the individuals' consent, insofar that the controller has initially collected the data by relying on a valid legal basis (as provided for by Articles 6 and 9 of the GDPR).

This means that, in case a hospital participating in the MHMD system intends to de-identify the data initially collected together with the patient's consent, with a view to carrying out further activities of non-personal data (such as for instance making those no-longer identifiable data available to third parties), that hospital should not go back to the individual to acquire his/her consent to put the anonymisation into effect. Rather, the application of pseudonymisation techniques does not require the data subjects' consent as such, but only on the basis of the further activity intended to be undertaken on pseudonymised data. Within the latter example, thence, the re-consent that the hospital acquire from the patient is not for pseudonymisation, but for communication of data to pre-identified third parties.

Interestingly for MHMD, homomorphic encryption (which is one of the most advanced encryption techniques explored in the project) can be considered as a valid form of pseudonymisation.

Inter-relation between data processing purposes and consent

In this case, the key-principle to be duly kept in mind is that each purpose of the processing requires a distinct and specific consent. This means, for instance, that should a controller wish to (i) carry out direct marketing campaigns, (ii) apply profiling algorithms to the customer base and (iii) communicate the data to third parties, it would have to obtain three separate and specific consents, one for each of the above-mentioned operations (insofar as no one of the alternative legal bases set out by articles 6 and 9 of the GDPR can apply in lieu of consent, at least for one or more of the processing).

Still, this general principle might be circumvented, as consent is for sure the main, but not the only legal basis for carrying out data processing activities. There are many other alternative legal grounds, mostly listed under art. 6 and 9 of the GDPR, such as when, *inter alia*:

1. the data subject has given explicit consent to the processing of those personal data for one or more specified purposes;
2. processing is necessary to protect the vital interests of the data subject or of another natural person where the data subject is physically or legally incapable of giving consent;
3. processing relates to personal data which are manifestly made public by the data subject;
4. processing is necessary for the establishment, exercise or defence of legal claims or whenever courts are acting in their judicial capacity;
5. processing is necessary for reasons of substantial public interest, on the basis of Union or Member State law which shall be proportionate to the aim pursued, respect the essence of the right to data protection and provide for suitable and specific measures to safeguard the fundamental rights and the interests of the data subject;
6. processing is necessary for scientific research purposes in accordance with Article 89.1 of the GDPR, based on Union or Member State law, which shall be proportionate to the aim pursued, respect the essence of the right to data protection and provide for suitable and specific measures to safeguard the fundamental rights and the interests of the data subject.

As evident, there are many cases when the data subject's consent is not strictly necessary. In these hypothesis, the controller would be required to prudently evaluate - on a case-by-case basis - whether any of the above listed legal grounds may apply to lawfully process the data.

MHMD smart contracts will take into account these principles, in order to be capable of checking the existence of any other legal ground (other than consent) to allow data access/processing, or to request a specific consent for different purposes pursued by the third-party data user.

Secondary use: when can this be lawfully done

Consistent with the fundamental principle of "purpose limitation", the data controller is required to evaluate with the utmost prudence whether any further processing of personal data which is intended to be undertaken may be considered compatible with the purpose(s) which has/have been made clear in the information notice initially provided to the data subjects.

As a general principle, any processing of personal data for purposes other than those for which the data have been originally collected shall be allowed, where no further legal basis exists apart from the consent originally provided by the individual, only where such processing is compatible with the original purpose.

In carrying out such compatibility evaluation, the data controller must take into account, in particular:

1. any existing link between the purposes for which the personal data have been collected and those associated to the intended further processing.
2. the context in which the personal data have been collected, with particular regard to the relationship in place between data subjects and the controller.
3. the nature of the personal data, in particular when falling within one or more of the special categories referred to by Art. 9 of the GDPR. In general, the more sensitive the information involved, the narrower is the scope for compatible use.
4. the consequences that may arise for data subjects from the intended further processing.
5. the implementation of appropriate safeguards, including encryption or pseudonymisation.

That being said, for the scope of MHMD, is also important to recall that Art. 5.1, let. b), of the GDPR, in defining the principle of purpose limitation, specifies that further processing for (*inter alia*) scientific research purposes shall be considered compatible with the initial purposes, where appropriate security measures and safeguards have been put in place pursuant to Art. 89.

Nonetheless, this provision must not be construed as an overall exception from the requirement of compatibility explained above, or a general authorization to further process data in all cases of scientific research. The Data Controller shall in any case carefully assess all the relevant circumstances and factors listed above to make appropriate decisions with regard to secondary data usage.

How to ensure the validity of consent (and consent method)?

There is no unique way of acquiring consensus. The GDPR establishes that it must be expressed in an *unambiguous* way and requires the controller to be able to prove its acquisition, when needed (which implies the adoption of ascertainable processes, such as the acquisition of a written signature or a tick in a specific online box or on written "Yes" or "Not" options).

Differently, as far as sensitive data (including health data) are concerned, the consent must be *explicit*, i.e., obtained in a totally demonstrable manner. At the end of the game, the output is very similar to that referred to personal data, since measures must be implemented in both cases which allow to prove by means of suitable documentation the acquisition of the individual's consent (e.g. by showing an IT evidence, a written form, or else). In case of sensitive data, however, it seems advisable and more prudent, to ask the data subject to indicate, both online and offline, his/her first name and surname, in order to make the consent as explicit as possible.

All these aspects will be taken in due consideration both for implementing the final dynamic consent form, and the relevant User Interface, providing that the evidence of the consent can be automatically checked (for authorising data processing) and easily retrievable (for *ex post* auditing procedures).

5.2.SMART CONTRACT LEGAL ISSUES

Although the potential of smart contracts is huge, as – once fully deployed – the technology could "decentralize the model of trust, speed up settlement times, reduce the need for costly intermediaries, enhance transparency, automate processes, reduce legal disputes, mitigate risk, and become the norm for countless types of transactions"[65], it is also true that various issues shall be addressed to make that potential become reality. This is particularly true whereas it is considered that "few lawyers have the coding skills to draft their own smart contracts [and thus] computer programmers would play a larger role, creating new liability questions for faulty algorithms and even ethical issues regarding the practice of law by non-lawyers"[66].

Indeed, the actual implementation of smart contracts can lead to various legal issues, in particular when the appropriate integration with the relevant legal prose is missing or insufficient. This is because the correct translation in code of the intended effect of the contract is a difficult task, and even small errors can lead to unexpected and huge undesired side-effect.

One of the most cited example of this is the DAO, an Ethereum-based decentralised autonomous organisation: the relevant smart contract contained a bug (known by the developers, who were actively working on a solution) which was subsequently exploited by one of DAO's participants to divert Ethereum (for an overall value of 50 million) to its own account.

[65] Getting Smart: Contracts on the Blockchain, Institute of International Finance, May 2016. https://www.iif.com/publication/research-note/getting-smart-contracts-blockchain.

[66] K. Silverberg, C. French, D. Ferenzy, Getting Smart: Contracts on the Blockchain, Institute of International Finance May 2016.

Two key potential questions/issues in translating the intentions in code can be identified[67]:

- "how do I know the code as written in the contract reflects my intentions if I cannot read it?". This question points to the already mentioned issue of making the smart contract understandable by non-technical readers (lawyers, party of a contract, etc.)

- "how do I know that the effect of the code, when executed by a machine, will be what I intend?". This question covers the possible risk of a bug in the code (as in the DAO example), leading to an unintended outcome.

To avoid unexpected outcomes, the following approaches/methodologies have been considered:

- Simulations: simulations of the smart contract behaviour shall be run, to check that the smart contract performs as expected. This could be done, in Hyperledger, through the "dev" mode[68], which can be used in the development phase for rapid code/build/run/debug cycles.
- Using standard/open source code: using standardised code – already tested and audited – for obtaining specific outcomes/elementary actions, could reduce the risk of unintended effects of the smart contract.
- Auditing procedures: it is becoming increasingly customary to have the smart contract "audited" by external bodies, as offered by OpenZeppelin with its Security Audit service[69], for ICO-related smart contracts.
- Working on contract automatic generation[70], with built-in "proof of robustness of the general code", starting with the preliminary identification of "erroneous coding patterns" and developing on top of that an automatic verifier for vetting chain codes for possible logical bugs.

In any case, as a part of the smart contract development framework, it is highly recommended to duly take into account the relevant liabilities linked to mistakes in programming, which include (but are not limited to) "product liability, breach of (the software as a service) contract, unfair and deceptive trade practices, and cybersecurity"[71].

Within D2.1, a preliminary list of best practices was provided, which is worthwhile recalling here briefly:

- *prepare a vulnerabilities memorandum*: due diligence should be performed before launching smart contracts, with a view to identifying potential vulnerabilities (involving legal, compliance and business personnel working with the smart contract developers to understand exactly what the self-executing and/or self-enforceable part of such contracts does (and doesn't)).

[67] As suggested in ISDA Smart Contract White Paper, op.cit.

[68] See relevant documentation: http://hyperledger-fabric.readthedocs.io/en/release/peer-chaincode-devmode.html.

[69] https://zeppelin.solutions/security-audits.

[70] S.Y. Chau, Making Blockchains Smarter: Toward Automate Robust Smart Contracts in Hyperledger.

[71] *Legal Aspects of Smart Contract Applications – Perkins & Coils White paper*, op. cit.

- *Take any reasonable measure to prevent risks*: assess any threat or risk that may reasonably occur and take any needed step to prevent it (proving adequate accountability).
- *Duty of cooperation*: sharing information with other participants to the Project and, in particular, to the smart contract, where feasible, will for sure help identifying and neutralizing vulnerabilities.
- *Contractually provide for contingencies*: the main remedy actions to be put in place following any smart contracts' breach, malfunction or unintended outcome should be pre-defined by the parties by means of specific agreements, to speed-up the recovery process and help mitigating the shortcomings.
- *Contractually prohibit consequential damages*: contractual terms intended to exactly outline the extent of reciprocal functions between the parties will make it much easier to allocate and eventually enforce the relevant liabilities.

Other specific issues can be identified in regard to smart contracts and GDPR

Besides considering the technical risks associated with the development of smart contracts, also their compliance with relevant regulation shall be considered.

As mentioned in D2.1, it is true that – in MHMD – one of the main functions of the smart contracts will be "matching the permission's extent pinpointed by each data subject by means of consents (as specified above, a specific consent must be acquired for each envisaged processing), with the access requests coming from the various stakeholders involved (research centres, businesses, etc.)", thus making the smart contract a means to ensure lawful data access and exchange.

At the same time, compliance with GDPR provisions shall be taken into account, in particular when the concept of "automated decision" is considered (as mentioned in recitals 15, 68, and in particular recital 71 of the GDPR, which states that "The data subject should have the right not to be subject to a decision […] based solely on automated processing and which produces legal effects concerning him or her […] without any human intervention").

As smart contracts can be surely considered "automated decision-making systems", subsequently art. 13, 14, 15, and art. 20 of the GDPR shall apply:

According to art. 13, and subsequent art. 14 and 15, whether or not the data has been obtained directly from the data subject or not, the data subject should be informed (and have access to the relevant information) about "the existence of automated decision-making", obtaining "meaningful information about the logic involved, as well as the significance and the envisaged consequences of such processing for the data subject".

This means that the existence of a smart contract-driven automated data exchange protocol shall be clearly explained to the data subject, in a form that makes it possible to the data subject to understand the underlying logic and possible consequences of this automated process.

Additionally, Section 4 of the GDPR, “Right to object and automated individual decision-making”, art. 22 (“Automated individual decision-making, including profiling”), states that “The data subject shall have the right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her”.

This right doesn’t apply if the automated decision is (par.2 of art. 22):

a) “necessary for entering into, or performance of, a contract between the data subject and a data controller;
b) authorised by Union or Member State law to which the controller is subject, and which also lays down suitable measures to safeguard the data subject's rights and freedoms and legitimate interests; or
c) based on the data subject's explicit consent”.

This means that within the consent form itself, an explicit consent option for automated processing shall be included and obtained from the data subject. In any case, as pointed out in par.3 of the same art. 22, “the data controller shall implement suitable measures to safeguard the data subject's rights and freedoms and legitimate interests, at least the right to obtain human intervention on the part of the controller, to express his or her point of view and to contest the decision”.

This last provision can lead to a technical implementation problem, whereas, once launched, a smart contract can’t be actually stopped. For this reason, in MHMD architecture, the idea of removing the link between a persistent identifier and the relevant dataset to which it points has been deemed to be a suitable solution for allowing the data subject to claim the right to be forgotten, or to re-consider his/her consent option before any other transaction involving his/her data would take place again. This solution, though, would not enable a data subject to reverse transactions already conducted on the basis of the previously granted consent (and with the original consent options), while future transactions can be held in abeyance until the explicit consent is not granted.