# Artificial intelligence Supporting CAncer Patients across Europe

| | |
|---|---|
| Project Title | Artificial intelligence Supporting CAncer Patients across Europe |
| Project Acronym | ASCAPE |
| Grant Agreement No | 875351 |
| Instrument | Research and Innovation action |
| Call / Topic | H2020-SC1-DTH-2019 / Big data and Artificial Intelligence for monitoring health status and quality of life after the cancer treatment |
| Start Date of Project | 01/01/2020 |
| Duration of Project | 36 months |

# D1.1 – Positioning ASCAPE's open AI infrastructure in the after cancer-care Iron Triangle of Health

| | |
|---|---|
| Work Package | WP1 – ASCAPE requirements, data structures, use cases and architecture specification |
| Lead Author (Org) | George Spanoudakis (STS) |
| Contributing Author(s) (Org) | Tzortzia Koutsouri (STS), Konstantina Koloutsou (STS), Miltiadis Kokkonidis (INTRA), Lucian Itu (SIE), Anamaria Vizitiu (SIE), Radu Toev (SIE), Tudor Suditu (SIE), Mirjana Ivanović (UNSPMF), Vladimir Kurbalija (UNSPMF), Miloš Savić (UNSPMF), Brankica Bratić (UNSPMF), Konstantinos Perakis (UBI), Dimitrios Miltiadou (UBI),Konstantinos Lampropoulos (UOP), Serge Autexier (DFKI), Johannes Rust (DFKI), Antonis Valachis (ORB), Manos Athanatos (FORTH), Thanos Kosmidis (CC), Paris Kosmidis (CC) |
| Due Date | 30.06.2020 |
| Actual Date of Submission | 30.06.2020 |
| Version | 1.0 |

**Dissemination Level**

X PU: Public (*on-line platform)

The work described in this document has been conducted within the project ASCAPE. This project has received funding from the European Union's Horizon 2020 (H2020) research and innovation programme under the Grant Agreement no 875351. This document does not represent the opinion of the European Union, and the European Union is not responsible for any use that might be made of such content.

☐ PP: Restricted to other programme participants (including the Commission)
☐ RE: Restricted to a group specified by the consortium (including the Commission)
☐ CO: Confidential, only for members of the consortium (including the Commission)

## Versioning and contribution history

| Version | Date | Author | Notes |
|---|---|---|---|
| 0.1 | 24.03.2020 | George Spanoudakis (STS) | TOC |
| 0.2 | 27.03.2020 | Miltiadis Kokkonidis (INTRA) | Initial revisions to TOC on the basis of input from T1.1 partners |
| 0.3 | 24.04.2020 | George Spanoudakis (STS) | First partner contributions merge and review |
| 0.4 | 12.5.2020 | George Spanoudakis (STS) | Second partner contributions merge and review |
| 0.5 | 25.5.2020 | George Spanoudakis (STS) | Third partner contributions merge and review |
| 0.6 | 15.6.2020 | Miltiadis Kokkonidis (INTRA) | Restructuring, minor corrections and review annotations of v3 contents on the basis of first internal review and subsequent discussions |
| 0.7 | 26.6.2020 | George Spanoudakis (STS) | Fourth partner contributions merge and review (focused on areas identified during the internal review) |
| 0.8 | 29.6.2020 | Miltiadis Kokkonidis (INTRA) | Corrections based on second internal review |
| 0.9 | 30.6.2020 | George Spanoudakis (STS) | Leading author's final check and corrections |
| 1.0 | 30.6.2020 | Miltiadis Kokkonidis (INTRA) | Work Package leader's final check and corrections |

# Table of Contents

## List of tables



## List of figures



## List of Requirements

## List of acronyms

| | |
|---|---|
| AHEE | Algebra Homomorphic Encryption scheme |

| | |
|---|---|
| AI | Artificial Intelligence |
| API | Application Programming Interface |
| ASR | Age-Standardised Rate |
| BDVA | Big Data Value Association |
| BFV | Brakerski/Fan-Vercauteren |
| BGV | Brakerski- Gentry-Vaikuntanathan |
| CKKS | Cheon- Kim-Kim-Song |
| CNN | Convolitional Neural Network |
| CPU | Central Processing Unit |
| CUDA | Compute Unified Device Architecture |
| DL | Deep Learning |
| DNN | Deep Neural Network |
| DP | Differential Privacy |
| EHR | Electronic Health Record |
| EU | European Union |
| FATE | Federated AI Technology Enabler |
| FedAvg | Federated Averaging |
| FedCS | Federated Client Selection |
| FHE | Fully Homomorphic Encryption |
| FIN | Fisher Network |
| GAM | Generalized Additive Models |
| GDP | Gross Domestic Product |
| GDPR | General Data Protection Regulation |
| GPGPU | General Purpose Graphics Processing Unit |
| GPU | Graphics Processing Unit |
| HDD | Hard Disk Drive |
| HE | Homomorphic Encryption |
| HELib | Homomorphic Encryption Library |
| HTTPS | HyperText Transfer Protocol Secure |
| ICT | Information and Communication Technology |
| ICU | Intensive Care Unit |
| IID | Independent and Identically Distributed |
| IT | Information Technology |
| LIME | Local Interpretable Model-agnostic Explanations |
| ML | Machine Learning |
| MLP | Multi-Layer-Perceptron |
| MORE | Matrix Operation for Randomization or Encryption |
| OpenCL | Open Computing Language |
| OS | Operating System |
| PHE | Partially Homomorphic Encryption |
| PII | Personally Identifiable Information |
| PINQ | Privacy INtegrated Queries |
| PRN | Partial Response Networks |

| | |
|---|---|
| PSA | Prostate Specific Antigen testing |
| QoL | Quality of Life |
| RFC | Request For Comments |
| SEAL | Simple Encrypted Arithmetic Library |
| SMPC | Secure Multi-Party Computation |
| SSD | Solid State Drive |
| SSL | Secure Sockets Layer |
| SVM | Support Vector Machine |
| TLS | Transport Layer Security |
| UI | User Interface |
| VMs | Virtual Machines |

# Executive Summary

There are over 3.7 million new cases of cancer in Europe and the number of people living with cancer is predicted to keep increasing. Motivated by the above, the aim of ASCAPE is to build an open Artificial Intelligence (AI) infrastructure for cancer patient support where valuable patient data-derived knowledge in the form of Deep Learning AI models from healthcare providers across can be collected and shared through the cloud while advanced technological means ensure patient data remain confidential. This data-derived knowledge is made available to doctors to aid them in their decisions and help provide a better Quality of Life trajectory to their patients. ASCAPE challenges the Iron Triangle of Health orthodoxy by offering opportunities for both Quality of Care and Access to Care to improve while the Cost of Care decreases.

This deliverable presents the methodology used for the process of gathering user requirements for ASCAPE Framework System. The results of this requirements specification are codified in generic outlines of functional requirements and a set of non-functional requirements.

In the ASCAPE workplan there are three tasks and two deliverables about specifications and requirements. The present deliverable focuses on the overall framework of ASCAPE, which is domain agnostic, whereas the follow-up deliverable D1.2 “ASCAPE Data Determinants and piloting validations”, due in Month 8, focuses on what aspects of quality of life are to be monitored for breast and prostate cancer, what will be the data collection framework in each pilot study and pilot-specific use cases.

The end-users of ASCAPE are categorised in three groups: doctors, patients and system administrators and for each group generic use cases are created. ASCAPE aims at providing a framework for creating ASCAPE-powered versions of existing software solutions used by healthcare providers and enhancing standard functionality in such systems. Functionality for user management, data entry, retrieval of patient records, etc. already exists in the Information System used by the healthcare provider. For this reason, only ASCAPE-specific functional requirements are presented. From the perspective of non-functional requirements, firstly we specify the security and the privacy requirements to avoid unauthorized access to patient data and ensure privacy preservation complying with the General Data Protection Regulation (GDPR) regulations. Then we investigate performance, hardware, usability, and overall quality requirements. The analysis phase identified nearly 50 requirements (including both functional and non-functional), that can be found assembled in the Appendix. In accordance with the Work Plan, requirements gathering focusing on pilot-specific requirements will be recorded in the follow-up deliverable D1.2 “ASCAPE Data Determinants and piloting validations” and an elaboration of requirements into a comprehensive system design and architecture for ASCAPE will be presented in D1.3 “Architecture Definition”.

# 1 Introduction

The latest cancer statistics show promising advances in decreasing mortality related to cancer (e.g. EU 4.5% between 2015 [1] and 2020, US 29% since 1991 to 2017 [2]). However, the number of patients living with cancer will grow significantly in the near future due to the fact that one in two people will be diagnosed with cancer in their lifetime [3], while at the same time the average life expectancy increases.

The main objective of ASCAPE is to take advantage of the recent Information and Communication Technology (ICT) progress in Explainable Artificial Intelligence, Federated Deep Learning and Privacy-preserving data processing techniques such as Homomorphic encryption and Differential Privacy, to improve cancer patients' quality of life and post treatment course. In order to reach its goal, ASCAPE will build an open AI infrastructure to enable health stakeholders (hospitals, research institutions, health care administrators/providers, etc.) to either deploy and run locally the AI algorithms on their private data, without sending them to the cloud, or sending their data homomorphically encrypted to the cloud where they may be processed without the cloud decrypting them. Even though patient data remain private, as they are not shared with external parties, new knowledge emerges from several AI analytics that is shared through the open AI infrastructure.

Within the context of ASCAPE project, two types of cancer are considered: breast cancer and prostate cancer. This way, coverage across genders, as well as age groups, will be achieved.

As Big Data Value Association (BDVA) states [4], only a radical breakthrough has the potential to disrupt the Iron Triangle of Health such that all three components including Quality, Access and Cost are simultaneously improved in cancer-care situations.

To this end, one aim of this deliverable is to carry out a set of detailed literature reviews and identify the advancements that ASCAPE will bring to the current state of the art and another to specify concrete requirements for the ASCAPE framework, whose realisation will enable us to deliver these advancements. Subsequently, this document is to define the use cases, identify functional and non-functional requirements, link them clearly to the use cases and prioritise them, thus laying the groundwork for the detailed system design in the forthcoming deliverable D1.3 "Architecture definition".

This document is structured in the following way. Section 2 provides the background of the project, explains the challenges that drive the need for ASCAPE solution and presents two use case scenarios as a mean of introducing the vision of ASCAPE. The state of the art is introduced in Section 3, focusing on explainable Artificial Intelligence (AI) algorithms for healthcare, federated deep learning, homomorphic encryption and epsilon-differential privacy. The advancements that ASCAPE brings to each of these technologies are explained. The requirement elicitation methodology adopted in ASCAPE is introduced in Section 4. Section 5 presents use cases for the three key stakeholders' groups: healthcare providers, patients and administrators. Section 6 presents the functional and non-functional system requirements for the ASCAPE framework. Section7 goes to the conclusions highlighting the most important points mentioned on this deliverable and the work to follow.

# 2 ASCAPE Value Proposition

## 2.1 Socioeconomical challenges landscape and ICT solutions

In the year 2018, the number of new cancer cases in EU was approximately 3.91 million (non-melanoma skin cancer was not included) while the number of deaths reached 1.93 million. According to [5] cases of female breast cancer (523,000), colorectal cancer (500,000), lung cancer (470,000) and prostate cancer (450,000) represent half of the total cancer cases in EU.
As far as breast cancer is concerned, according to the CONCORD-3 study and based on data from 2010 to 2014, the 5-year net survival age-adjusted probability in all adults, in the 28 countries of the European Union (EU), ranges from 79% in Croatia to 93% in Cyprus [6]. In 2018, the 5-year prevalence (number of people who have had a cancer diagnosis in the last 5 years) for breast cancer was in absolute number of 2,054,887, from a total of 12,132,287 total cancer prevalence [7].
Regarding prostate cancer, the approximate number of new cases in EU in 2015 is about 365,000 and is the most frequently diagnosed type of cancer in men. The incidence rates (ASR: age-adjusted rate on the European standard population) in EU range from ASR 175 in Sweden to ASR 34 in Greece. The 5-year prevalence of prostate cancer in EU is about 1,300,000, while at the same time survival has raised in all the EU countries with the highest increase monitored in the Eastern countries. The introduction and wide use of Prostate Specific Antigen (PSA) testing and diagnostic procedures such as biopsy, have changed the distribution of the disease [8].
According to the numbers reported before, breast and prostate cancer survivorship represent a huge health problem for European countries. Breast and prostate cancer patients present psycho-social needs. Physical, social and emotional scars could compromise return to everyday life. Different studies showed that almost a third of cancer survivors experienced changes in their work situation after treatment [9]. Some of the most common problems in returning to normal life after cancer is obtaining life or health insurance and home loans. The patient-centred approach is fundamental for improving the Quality of Life (QoL) of cancer patients through rehabilitation and support. However, two main obstacles restrain health providers from using survivorship care plans often: 1) the feasibility of integration of distinct health levels and 2) the cost and resources required to develop and manage these plans [10].
Socioeconomical challenges of monitoring health status and QoL after cancer treatment involve resources, demography, economy, society, and governance.
The resources to address those challenges could be diverse, but undoubtedly the most used in health and medical care are smartphones and wearables.

- At the end of 2017, 85% of the population of EU was using a mobile connection that results in 465 million unique mobile subscribers. At the end of 2019, the percentage changed to 86%, while 76% of total connections has adopted a smartphone. In the year 2018, mobile services and technologies produced 3.3% of European GPD – a value-add of €550 billion. By 2022, this number is estimated to raise to €720 billion (4.1% of GDP) [11].
- The global market in 2019 reached 3,800 million mobile internet users and

8,000 million SIM connections [12].
- The number of wearables in 2017 was estimated at 527 million, 116 million from those in Europe [13].

More individuals are monitoring their health through wearable devices, such as fitness trackers. However, monitoring diseases like cancer, still faces some barriers, for example, the interoperability with health centre informatics systems and electronic health records (EHRs), or the extraction of the value of the information produced by wearables (directly through a Bluetooth or WiFi connector or via vendor's data cloud).
Demographics also matter in the use of smartphones and wearables. In 2018, nearly one fifth (19%) of the EU population was aged 65 and more [14] [10]. There still exists a digital gap between older and younger generations. The main factors are generation effects, cognitive and physical deterioration connected to aging, and unfavourable approach towards technology. Nevertheless, the number of older adults that adopt different kinds of technologies to fit in with society is increasing [15].
Another aspect that can hinder access to technologies is economic. In 2017, 6.6% of the population in the 28 countries of the European Union was extremely materially deprived. The highest percentage of individuals being at the risk of social exclusion and penury was recorded in Bulgaria (38.9%), Romania (35.7%) and Greece (34.8%). On the other hand, the lowest one was monitored in Slovakia (16.3%), Finland (15.7%), and Czech Republic (12.2%) [16].
Access to education, access to health system, as well as the quality of living conditions, are important factors to addressing the challenges of monitoring QoL after cancer treatment. Social perception, efforts devoted to information and communication, training, qualification and participation are important to achieve better use of ICT for monitoring health status.
As far as governance is concerned, the EU regulation on data protection [17] has a special impact and created both new challenges and opportunities for advanced technological solutions such as ASCAPE.
The ASCAPE platform will have to adhere to GDPR, to national data protection, privacy and ethical legislation and provide an effective means of assisting doctors to help patients navigate the challenges of their disease, the side-effects of treatments and a host of other factors in the context of a diverse range of socioeconomic challenges within which it will need to operate and prove itself to be an effective ICT innovation with a positive social impact.

## 2.2 Addressing the challenges of the Iron Triangle of Health

In economics terminology, demand is always associated with the customer's readiness to pay for the desired product or service. However, in healthcare economics the idiosyncracies of the demands and the complexity of the services require a broader interpretation as other factors arise, like the right to access the healthcare services, and the demands are always greater than the resources. That is reflected in the paradigm of "the Iron Triangle of Health Care" coined by William Kissisk in his book Medicine's Dilemmas: Infinite Needs Versus Finite Resources (1994) [18]. The "Iron Triangle of Health" represents the concept of three opposing features in the healthcare system: healthcare access, delivery quality and cost of the care. The underlying idea is to demonstrate the interconnections among the three mentioned features, stating that any impact on one of them will also impact the other two. Therefore, it would be almost impossible to improve access and quality while simultaneously reducing costs. For example, making healthcare accessible to more people without increasing costs necessarily needs to impact negatively the quality of care, or the quality of care could be improved but it will require the increment of cost or the limitation of access to healthcare. The concept is widely used in health policy and health economics with the general goal of keeping the balance under the inalterable constrain that increasing one feature will impact the other two assuming that there are not unnecessary services, inefficient processes and unfair prices.

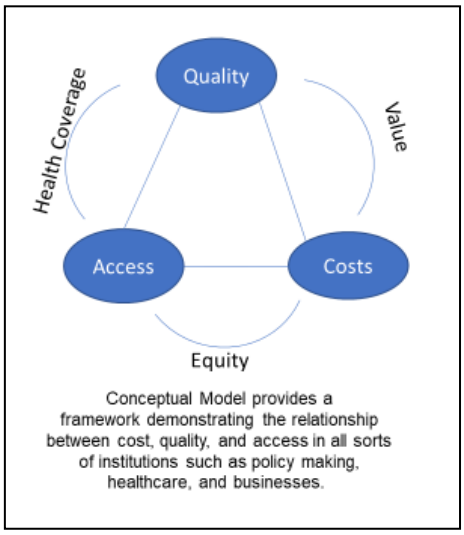


Figure 1. Health Iron Triangle model. Access means how easily can patients get healthcare services; quality refers to how good are the services delivered and cost refers to how expensive is the service

The Iron Triangle concept is not written in stone but could help us understand and improve aspects of healthcare. Healthcare systems face the lack of human and economic resources and the ever-increasing demand for quality services. Still, the Iron Triangle concept is critiqued as proposing a rigid model that does not take into account costs change dynamics. Another major criticism is related to the impact of ICT in healthcare, and to the fact that the application of technology has already disrupted the model. The current development of technology, and AI in particular has already challenged the Iron Triangle concept and holds the promise that health care access and quality could be improved and cost reduced. Certainly, the development of AI has impacted the industry, socio-economics interactions, biomedical research as well as healthcare provision. AI is able to improve the quality of care by supporting better diagnosis, treatment and care. Moreover, AI can extend healthcare coverage by optimizing the resources and at the same time reducing cost. According to Forbes publication [19], the investment in AI in the healthcare sector will reach $6.6 billion by next year. The same publication also refers to an Accenture report that estimates that applying AI may result in $150 billion saving by 2026.

ASCAPE aims at disrupting the Iron-Triangle paradigm by developing AI models that will improve the Quality of Life (QoL) of cancer survivors while reducing costs to the healthcare systems and improving access to services.
The opportunities for disrupting the Iron Triangle paradigm stem, among other factors from:

- ASCAPE's ability to collect input directly from cancer patients and their devices (online questionnaires, mobile apps, wearables) which translates to:
  - more data (which could lead to increased AI results and hence to better Quality of Care)
  - less administrative effort for collecting the data (leading to lower personnel costs i.e. lower Cost of Care, as well as to freeing time and human resources that could be re-allocated to improving Access to Care)
- ASCAPE's envisaged ever-increasing quality of AI support offered to doctors, which could:
  - lead to doctors making better recommendations to their patients, leading to better Quality of Care, which may, in turn, result in lower Cost of Care (as better recommendations may also be more cost effective straight out or in the long run)
  - predict patient quality of life issues and notify the doctor early leading to better Quality of Care and possibly lower Cost of Care at the same time (as the early identification of impaired quality of life issues would enable an earlier, less evasive, and probably more cost-effective treatment)

The forthcoming deliverable D1.4 "Manuscript on costs and benefits of the new diagnostic tool" will provide a theoretical insight into the cost-effectiveness model for ASCAPE aiming at providing a better understanding of the potential AI methods for improving breast and prostate cancer patients QoL.

## 2.3 ASCAPE: Changing Healthcare for Patients

ASCAPE is a complex technological undertaking with a bottom line measured in a variety of ways, one being patient experience of an ASCAPE-powered healthcare system. The following two scenarios present the ASCAPE vision from that angle.

### 2.3.1 Use Scenario 1 (Peter, prostate cancer patient)

Peter is invited to update his health status on the ASCAPE Patient service. He responds to a series of questions that address his emotional wellbeing, his direct physical side-effects, the state of his intimate relationships, and more. Among these questions is one that is quite relevant to his recent concerns: "Have you been constipated?"
He records his response without hesitation: "Very much".
He had been looking forward to his doctor's views on this very frustrating experience, which he believes is a side-effect of his chemotherapy; consequently, he was planning on asking during his next appointment, a couple of days later.
To his surprise, his doctor had been prepared to discuss this as well: she had seen his input and had ran the ASCAPE algorithms to review the available information and

options. She discussed the two top options with him, which were either to take a mild over-the-counter medication, or to engage in a medium-intensity exercise plan. She answered his questions and based on their discussions, she recommended to go with the exercise plan, so as to avoid polypharmacy.
Beyond constipation, however, his doctor had also identified a concerning sign: his answers on emotional wellbeing were hinting at a slight deterioration. When she asked him about this, Peter shrugged it off, saying that "he's fine" but she insisted that the ASCAPE predicted a deterioration for people like him. She suggested that he remains aware and conscious of his wellbeing, for now, and made sure to comfort him regarding how common and manageable this is, if diagnosed early.
Later, and as Peter was engaging with the ASCAPE Patient service again, he was happy to report that his constipation had subsided from "Very much" to "A Little", possibly thanks to the exercise he had decided to stick to. Interestingly, he was looking forward to the questions on emotional wellbeing; he took some time to respond to them as he remembered his doctor's words. He indeed, felt OK – but also felt reassured by the fact that he can remain in control if he and his doctors can read the signs on time.

**ASCAPE impact**

In this scenario, ASCAPE enhanced the healthcare professional's interaction with the patient by collecting (and analysing) quality of life data and indicating possible interventions, which she used to make a recommendation to the patient. This, in turn, can improve the patient's quality of life.
Moreover, the prediction about the patient's emotional deterioration helped (i) identify an important aspect of his quality of life, that may have gone unnoticed, and (ii) make a prediction using incomplete data. This sets the foundation for improved awareness of health status, and improved quality of life.

### 2.3.2 Use Scenario 2 (Alice, breast cancer patient)

Alice has been patiently entering her information on the ASCAPE Patient Service. Her journey with breast cancer has been relatively smooth, with minor side effects that generally come and go but do not affect her too much.
Recently, however, something strange has started to happen: every time she picks up the pen and starts to write, her handwriting looks unfamiliar to her. She feels that the grip she has on the pen has changed, and the feeling of the pen itself on her fingers is strange.
The next time she enters her information, she notices a question about whether she feels "numbness" in the fingers/toes. She decides that this is the most relevant question to her experience, and after deliberating with herself over the various options for the response, she picks "Quite a bit". She moves away from her computer, picks a blank piece of paper and a pen, and starts to write. The feeling confirms her response – indeed, her fingers feel numb.
During her next appointment with her doctor, he opens up her digital files and notices the difference from the previous inputs. He discusses more about this with Alice and introduces her to the concept of peripheral neuropathy. He answers her questions and then pulls up the ASCAPE recommendations for intervention. These include multivitamins, omega-3 fatty acid supplements, acupuncture, and yoga. The

recommendations include some additional details that help the doctor make the final suggestion to Alice.
His suggestion is for Alice to start taking multivitamins in order to control the symptoms. Alice agrees, but asks about “Plan B”: what if this does not work? Her doctor is understanding of her concerns and reassures her that they can reconsider the options together in a few weeks.
He also pulls up some Prediction metrics from ASCAPE which he feels is reasonable to share with Alice. Based on these predictions, patients taking the medications that Alice has been taking have a medium to high risk of experiencing these symptoms – so this is not uncommon.
Alice would wish that her symptoms magically go away, but recognises that this is not so easy, and her treatment is very important to her. At the same time, she feels more confident that she stays in control, and even has a “Plan B” with her doctor, if necessary.

### ASCAPE impact

In this scenario, ASCAPE helped the patient identify and report one of her side-effects that is hard to pinpoint and track. This, in turn, supported the healthcare professional’s diagnosis of the side effect and enabled him to discuss with the patient about possible intervention options. As a result, the first step for managing this side-effect was made. However, these intervention options are not always without risks. ASCAPE made the conversation between the patient and her clinician much clearer, by helping predict what could be the consequences of the intervention options presented (based on predictive models). As a result, this allows the patient to stay in control and know what to look out for, should the need for a change in plan arise.

# 3 ASCAPE Technologies State of the Art and Beyond

ASCAPE aims to advance the state of the art in healthcare applications of four areas of AI research: explainable AI, federated deep learning, homomorphic encryption-supporting AI and privacy aware-AI based on epsilon differential privacy. For each of these four areas the relevant state of the art is presented together with concrete proposed advancements with respect to the status quo. The proposed advancements constitute initial research directions, based on the project needs and the Consortium's prior experience, which ASCAPE may follow, adjust, and move beyond as ASCAPE-specific experience is gained, problems are identified and solved, and alternative methods are explored. They also provide a bold early statement of the ASCAPE Consortium's plan of attack on the practical challenges of the project and as such relate to meeting specific requirements as elaborated further in Section 6. ASCAPE state of the art advancements are likely to be scientifically interesting individually, but also play a role in building a technologically disruptive healthcare AI infrastructure which has the potential to transform the reality of cancer patient healthcare in the foreseeable future.

## 3.1 Explainable AI for healthcare

### 3.1.1 Current State of the Art

AI has a long tradition in computer science dating back to end of the 1950ies and has recently regained interest due to the practical success of Machine Learning (ML) techniques. Learning always was a central topic in order to build usable AI-systems that can learn from data, extract and generalize knowledge and extract factors that underlie and can serve to explain the data. The selection of appropriate features and the quality of the data and their interpretation taking into account the application domain provides the best results [20] and at the same time is a major challenge.
The medical domain has been one application domain for AI systems since the early 1970s (e.g., [21]) to support the detection of drug interaction, medical diagnostics and decision making. AI explainability is the goal of ensuring AI can present its results in a manner that makes sense and creates confidence to human users, rather than having the AI operating like a magic black box producing arbitrary answers. Explainability in the medical domain is required in order to enable medical professionals to understand why and how a decision is made in order to (gain) trust in the systems. The ability to provide explanations is the basis to design a suitable user interface to interact with medical professionals, which is a topic for Human-Machine-Interaction-Design and depending on the appropriate interaction-style imposes usability and operational requirements to the explanation system [22].
The early medical support systems such as [21] were mostly built on rules close to formal logical reasoning rules encoding knowledge of the domain experts. Knowledge was represented symbolically and the inference engine applying the rules could provide the used rules as a basis for explanations.
Medical decision support increasingly relies on ML models due to their flexibility and predictive power based on data and with impressive results even without domain

knowledge expertise. However, applications in routine clinical care are scarce. This is partly because trust in the safety of these algorithms is lacking, and because acceptance of complex models requires user interfaces that are readily understood by clinicians. Depending on the ML method, explanations are clear from the method (e.g. decision trees [23] have this advantage, though they do not perform as well as other methods) or inherently difficult to obtain (e.g. deep learning is known for its good results but explaining how they were reached is not easy). The development of explainable AI-systems typically has to resolve resp. balance between difficult to explain but highly performant ML methods and good to explain but less performant ML methods.

ML model interpretability has been addressed in the past for Multi-Layer-Perceptron (MLP) [24] [25], Support Vector Machine (SVM) [26] [27], Fuzzy logic [28] [29] and deep neural networks [30] [31]. Approaches to explain ML algorithms currently fall into four broad categories:

1. **Feature attribution**: attributing the classification to a small number of numeric/ semantic features. These feature attribution methods are usually interpretable by design. However, it is difficult to derive their form from data in a computationally efficient manner. Some advances have been made in generating nomograms for flexible models applied to tabular data [26] [27] [25] [31].
2. **Saliency maps**: sparse components of the original signal are identified, that have most influence on the model predictions, e.g. LIME [32].
3. **Activation maximization,** for example, based on Generative Adversarial Networks [33], which allows to determine which inputs maximize confidence in the output;
4. **Metric learning**: consists of deriving a metric from a classifier and using it to map out the data structure [34]. Similarity networks are generated from which a classification of an individual case can be obtained by consulting its neighbours. Additionally, explicit Siamese Networks have become very popular recently [35].

Approaches to build explainable prediction systems can be broadly classified as ante-hoc (using models that are interpretable by design) and post-hoc (methods interpreting black-box models).

***Ante-hoc explainable methods***

These models follow the principle that the best explanation of a simple model is the model itself; it perfectly represents itself and is easy to understand. This refers to making a complex model, e.g. MLP, interpretable by design. These methods aim at obtaining a model of the whole model. Typical examples are linear regression and decision trees.

1. **Partial response networks**: Partial response networks (PRNs, [24] [25]) are a novel representation of MLP. It is a constructive framework that explicitly models the output from clear and transparent features. Each feature depends on a small number of variables and has an attributed weight that is readily understood by the user. The contributions of the features are added together in a generalized additive models (GAM) to derive the probabilistic inference of

class membership. Importantly, the full model is derived from the original MLP and trained again, and Partial Response Networks (PRN) performance is comparable to that of traditional MLP and other related approaches. Additionally, it is able to discover the features that contribute meaningfully to the output, and to show how the output is built-up from them.

2. **Evolutionary Fuzzy Modelling**: Fuzzy logic systems can make accurate predictions, while providing a reasonable level of interpretability [28] [29]. It is based on a modelling approach capable of automatically building and testing its own set of rules with the help of an evolutionary algorithm. This approach was successfully applied in many contexts, including biomarker discovery and cancer diagnosis, leading to a commercial solution allowing for the discovery of interpretable diagnostic signatures [36].

**Post-hoc explainable methods**

Rather than trying to obtain a whole model of a system, these models aim at obtaining explanations of individual decisions from models treated as black-boxes.

1. **Activation maximization**: Using deep generative networks and tailored optimization methods, this approach automatically generates class-relevant images for any trained Convolutional Neural Network (CNN) [30]. A human user can then understand the internal representations assimilated by the network and the typical representations of the classes (representations may be easily labelled by human experts).
2. **Rule extraction** based on decision trees: This method allows for the extraction of decision rules from deep neural networks to transfer knowledge from a reference model into an explainable equivalent [31]. It consists of three main steps: (i) a trained CNN extracts features from a dataset, (ii) a Random Forest is trained to create rules based on such features, and (iii) the rules are ranked and selected according to their contribution conserving prediction performance while adding explainability.
3. **Fisher Networks** [34] is a framework that extracts from a dataset with indicator labels, an interpretable representation in the form of a similarity network informed by a given query about binary or multiclass assignment. The underlying structure of the network reflects the statistical geometry of the original data space as determined by density function estimates. It is then straightforward to visualize similarity even for high dimensional data. As a post-hoc method, FINs are typically applied to the final dense layers of trained deep networks.

### 3.1.2 ASCAPE Beyond State of the Art

A recent work [37] proposes a systematic assessment of explainability approaches to compare the different approaches developed thus far and derives a fact sheet describing the explainability methods. Taking the perspective from the learning task to explain in a specific domain and for different stakeholder explainees, they define five

dimensions of requirements an explainability method needs to fulfill in order to be adequate. The dimensions are:

- **Dim. 1** Functional requirements specifying (e.g. supervision level, problem type, explanation target, explanation scope, etc.).
- **Dim. 2** Operational requirements characterising how the stakeholders interact with the explanation system.
- **Dim. 3** Usability requirements characterizing properties of explanations that are important to the explainees.
- **Dim. 4** Safety requirements characterizing the effect of explainability on robustness, security and privacy aspects of predictive systems**.**
- **Dim. 5** Validation requirements specifying how the explanation system itself should be validated.

ASCAPE developing an AI-support for predicting and improving the QoL of cancer patients to be used by medical persona, we propose to apply a user-centered-design approach to develop the explainability system for the ASCAPE models. The assessment dimensions used in the fact sheet [37] can serve as a starting point to derive the requirements from application context and stakeholders requiring explanations in order to adopt the AI-support in practice. From these requirements the research and development process for ML methods to use as well as security and privacy relevant design decisions are not only informed by the properties of the learning methods and legal regulations, but also by the needs of the target stakeholders.

Providing the explainability capabilities in the ASCAPE platform will be a process in the field of tension created by on one hand side the requirements along the five dimensions identified by the target stakeholders and application context and on the other hand the need for models with high predictive quality and possibilities to obtain explanations depending on the ML algorithms used to train the models. The advancement beyond the state of the art of explainable AI will thus be driven by these two factors.

#### *3.1.2.1 Advancement 1 – Determining key determinants for impacts on QoL*

The classical role of explainability is to serve as tool to provide insight to the decision process of the model. In ASCAPE we will investigate how to use the possibility to reverse the information flow in the model by mapping the output to the input variables, it can also help to identify data points in the patients' data which have impact on the patients' quality of life. Explainability can also assess the significance of a data point and thus evaluate a proposed intervention regarding its expected effectiveness.

#### *3.1.2.2 Advancement 2 – Using explainability to predict model outputs*

Based on Advancement 1, ASCAPE will investigate how explainability can be alleviated to have a predictive capability. For instance, in combination with Monte-Carlo-Methods, explainability can be used to automatically identify medical interventions and propose these to the medical staff.

#### *3.1.2.3 Advancement 3 – Combing explainability on machine learning with human knowledge*

In an application domain agnostic setting, explainability provides information on the influence of input variables to decisions. However, in the medical context with its planned use for risk assessment as well as for intervention suggestions, input variables that can be influenced by interventions are more informative than those, that cannot be changed. For instance, gender or age of a patient are less informative compared to physical fitness or lifestyle habits regarding the identification of targets of interventions. In ASCAPE we will research how to feed that knowledge into the methods to generate explanations, that can be better used by the medical professionals.

## 3.2 Federated deep learning for healthcare

### 3.2.1 Current State of the Art

Federated learning is a ML technique which enables the use of decentralized data, e.g. residing on devices [38]. There is an increasing amount of data produced by healthcare organizations worldwide providing both advantages and challenges [39]. ML provides the tools needed to analyse big data. Federated learning attempts to solve the data dilemma faced by traditional ML methods by enabling the possibility to train a shared global model with a central server, while keeping all the sensitive data in local institutions like hospitals [40]. As the title suggests, the purpose of this section is to analyse the state of the art on federated learning in healthcare.
**Federated learning challenges.** The main challenges that arise from federated learning algorithms are [41]:

- **Statistical:** Any data points available locally are far from being a representative sample of the overall distribution.
- **Communication:** The number of clients may be large and can be much larger than the average number of training samples stored in the activated clients.
- **Privacy and security:** It is impossible to assume that none of the clients are malicious.

We will briefly describe the challenges enumerated above.
**Statistical challenges of federated learning**. Federated Averaging (FedAvg) was proposed to solve the federated learning problem but the performance of convolutional neural networks can significantly drop due to weight divergence [42], [43]. Another problem for FedAvg is that it does not address the statistical challenge of strongly skewed data. The authors of [40] classified the existing statistical challenges of federated learning in two groups: consensus solution and pluralistic solution. For consensus solutions a proposed solution is to model the target distribution or force the data to adapt to the uniform distribution [43], [44]. Another method that is employed is the sharing of a small portion of the data. Many researchers choose pluralistic solutions because it is hard to find a consensus solution that is good for all components. Corinzia et al. [45] introduced VIRTUAL, an algorithm for federated multi-task learning with non-convex models.

**Communication challenges of federated learning**. The training data is distributed over a large number of clients. Most of the times the clients have unreliable and relatively slow internet connection. The main challenge is to make the communication efficient and to reduce the data exchange between the clients and the server. In federated learning, there are three ways to solve this issue: reduce the number of clients, reduce the update size and reduce the number of updates. Based on the three points enumerated previously the authors of [40] classified the existing research for federated learning communication efficiency into four groups: client selection, model compression, updates reducing and peer-to-peer learning.

- **Client Selection.** Client selection is based on restricting the participating clients or choosing a fraction of parameters to be updated at each round. Protocols like the selective stochastic gradient descent protocol [46] or FedCS [47] are meant to solve the client selection problem.
- **Model compression.** The goal of reducing the communication cost is to compress the server-to-client exchange. A multi-objective federated learning was proposed by authors in [48], to maximize the learning performance and minimize the communication cost.
- **Updates reducing.** Kamp et al. [49] proposed to average models dynamically depending on the utility of the communication. This is well suited for massively distributed systems with limited communication infrastructure.
- **Peer-to-Peer learning**. In federated learning a central server is required to coordinate the training process of the global model. To solve this problem, Roy et al [50] proposed BrainTorrent, where all clients can interact with each other, without depending on a central body.

**Privacy and security challenges for federated learning**. In federated learning, the number of participating clients is large (up to one million). We have to take into consideration that some of these clients can be malicious. Yang et al. [51] introduced a comprehensive secure federated learning framework. Some researchers explored and indicated the vulnerability of the federated learning setting [52], thus accelerating the need of an effective defense strategy.
Among the defense strategies we can enumerate: Secure Multi-Party Computation (SMPC) [52] or differential privacy (DP) [53]. SMPC cannot prevent an adversary from learning individual information, and DP only protects users from data leakage to a certain extent. Another problem is that SMPC protocols are computationally expensive, even for the simplest problems and DP may reduce performance, in terms of prediction accuracy, Truex et al [54] combine DP with SMPC to reduce the growth of noise as the number of parties increases, without sacrificing privacy. The current utility protocols for privacy and security work only if the server follows the protocol.
**Existing applications in healthcare with federated learning.** Applications of federated learning in healthcare can be found in different scenarios: linear regression, logistic regression, object detection, or image segmentation. Some researchers used federated learning in linear regression problems to predict future hospitalizations [55], mortality rate and hospital stay time [56]. Huang et al. sought to tackle the challenge of non-IID (Independent and Identically Distributed) ICU (Intensive Care Unit) patient data, that complicated decentralized learning, by clustering patients into clinically

meaningful communities, and optimizing the performance of predicting mortality and ICU stay time. Brisimi et al. [55] aimed at predicting future hospitalizations for patients with heart-related diseases, using EHR data spread among different data sources, by solving the regularized sparse Support Vector Machine classifier in federated learning environment. Authors from [57] proposed a federated learning system for brain tumour segmentation and studied various practical aspects of the federated model sharing while preserving patient data privacy.

**Existing frameworks for federated learning.** Given the growing popularity of federated learning, several companies and research teams developed federated learning frameworks:

1. **Tensorflow federated.** Tensorflow Federated is an open source platform for ML and other calculations to be performed on distributed data [58]. It was originally developed by researchers and engineers working on the Google Brain team within Google's Machine Intelligence Research organization, to conduct ML and deep neural networks research. The major advantage of Tensorflow Federated is that it has a large and active community. The major disadvantage is that it is quite difficult to learn and to debug [59].
2. **PySyft.** PySyft is a library for implementing federated learning from the open-source community OpenMined [60]. It enables secured, private computations in deep learning models. The principles of PySyft were originally outlined in a research paper, and later on, it was implemented by OpenMined, which is one of the leading decentralized AI platforms [61]. Amain advantage in using PySyft is the wide range of tutorials they offer, including a free course on Udacity, that teaches users federated learning in PySyft.
3. **Substra.** Substra is another framework that is used for federated learning. Substra development started in April 2018 in Nantes (Frances) by Owkin's Substra team [62]. In 2019 it engaged in large collaborative research projects in Europe, in the health sector. The main contributor for Substra is Owkin, which is a fast-growing health data AI startup. Substra is a framework for traceable ML orchestration on decentralized sensitive data. Unlike Tensorflow Federated and PySyft, Substra is not that well documented, the only place where one can find tutorials is on their GitHub repository.
4. **FATE.** Federated AI Technology Enabler (FATE) is an open-source project initiated by Webank's AI department to provide a secure computing framework to support the federated AI ecosystem. The protocols that it implements are based on homomorphic encryption and multi-party computation [63]. FATE can only be installed on Linux or Mac, representing a disadvantage for those who rely on Windows as their main operating system (OS).

In this section we described the state of the art for federated learning in healthcare, we enumerated some algorithms used in federated learning, what applications were made in healthcare with federated learning and some frameworks that can be used for future researches in this area. Some future research for federated learning in healthcare should be done for improving data quality, incorporating expert knowledge and improving model precision in federated setup.

### 3.2.2 ASCAPE Beyond State of the Art

From the overview of the current State of the Art given in the previous section, it can be observed that existing federated learning methods and frameworks for healthcare applications consider federated learning in controlled settings in which data federation partners (clinics in our case) must be completely known in advance. In such controlled settings, the learning of a federated model is driven by a federated learning server which waits for all federated clients to connect in order to start the learning process. In the first learning round, all federated clients train individual models on their local (private) data and send the trained models to the federated learning server. The federated learning server then reduces (averages) the received models (e.g. by averaging link weights of neural networks having the same structure) and sends the averaged model back to all federated learning clients or to a subset of them according to its own client selection policies. The federated learning clients update the parameters of the averaged model on their local data and return the updated models back to the federated learning server for the next round of averaging. The whole process is repeated for an arbitrary number of learning rounds. This federated learning scheme can be named as the concurrent federated learning since data federation partners update the global model in parallel being carefully synchronized by the federated learning server.

#### *3.2.2.1 Advancement 1 - Incremental and semi-concurrent federated learning schemes for cancer-care predictions*

The primary purpose of federated learning in ASCAPE is to enable democratized access to machine learning models promoting cancer patient quality of life without revealing private or sensitive patient data. It is evident that the concurrent federated learning is an inadequate approach for democratized federated models since: (1) data federation partners (clinics) are not known in advance, and (2) data federation partners may constantly change in time. Thus, in ASCAPE we will propose a novel federated learning scheme in which federated models are learned incrementally or semi-concurrently as clinics join to the ASCAPE platform. In the simplest incremental learning setting, a federated learning client joining the ASCAPE platform downloads the global model from the ASCAPE cloud, updates the parameters of the global model with its local data on the local ASCAPE edge node and sends the updated model back to the ASCAPE cloud. The first registered data federation partner trains the initial model. Additionally, the global model is locked while it is being updated by a data federation partner on its local ASCAPE edge node. This means that other data federation partners wanting to update the global model wait for the partner currently updating the model to finish with the update process (meanwhile those partners waiting to update the model may use the old global model to make predictions or predictions may be made by a personalized model trained only on local data; this personalized model is trained in any case as we will explain later). After a data federation partner updates the model it is redistributed to all registered clinics (i.e, sent to their ASCAPE edge nodes).
We will also consider semi-concurrent federated learning schemes in which multiple data federation partners joining the ASCAPE platform at close time intervals may

update the model concurrently. Additionally, we will investigate model rollback options taking into account that there will be data federation partners having ASCAPE edge nodes preserving local data (without sending the data to the ASCAPE cloud). With model rollback options semi-concurrent federated learning can be less sensitive to non-IID patient data. In any case, the federated learning schemes designed in ASCAPE will assume that data federation partners are not known in advance and that they may join at any time, thus enabling democratized access to knowledge captured by trained ML models.

#### *3.2.2.2 Advancement 2 - Personalized cancer-care predictive models in federated learning settings*

For each ASCAPE outcome variable (e.g., a QoL indicator or an intervention) there will be a dedicated global model stored in the ASCAPE cloud. When a clinic joins the ASCAPE platform, the global model for a particular outcome variable will be downloaded to the local ASCAPE edge node and evaluated on local training data. In the case that the global model exhibits a high accuracy then it will be updated on local data and sent back to the ASCAPE cloud. Otherwise, a personalized model for the clinic will be trained considering only local data. All personalized models will be also stored in the ASCAPE cloud and together with the global model they will constitute an ensemble of models for a particular outcome variable. With subsequent clinics joining the ASCAPE platform all models from an ensemble will be evaluated on the local data and the best one will be selected for update and later use when making predictions. Additionally, we will consider prioritizing the models from the ensemble according to the number of updates and previously observed accuracy values. With a significant number of models in the ensemble ASCAPE will also enable informed ensemble-based predictions.
ASCAPE federated learning will be augmented with feature selection techniques. Both filter and wrapper methods will be utilized to identify a subset of features leading to the most accurate personalized models. Additionally, we will also explore recently proposed graph-based methods based on community detection algorithms applied to feature correlation networks [64]. The accuracy of global models retrieved from the ASCAPE cloud will be also evaluated in their reduced forms (e.g. dropout in neural networks) after feature selection on local training data in order to see whether reduced global models can yield more accurate predictions. This will be the second way for making personalized models in ASCAPE.

#### *3.2.2.3 Advancement 3 - Semi-supervised cancer-care predictive models in federated learning settings*

Existing federated learning algorithms mostly focus on supervised learning of classification and regression models. Besides supervised learning, the ASCAPE platform will also support semi-supervised and unsupervised learning of federated machine and deep learning models combined with feature selection techniques. Regarding semi-supervised learning scenarios, personalized models trained on local data will be used to infer missing values for outcome variables. This data inference

process will be realized incrementally considering both the personalized model to infer missing values and the global model to check predictions:

1. An initial version of the personalized model is trained from existing labelled data (data instances with known values of the outcome variable).
2. The personalized model is applied to unlabelled data to obtain predictions for missing values of the outcome variable.
3. Predictions made by the personalized model are checked by the global model. Those data instances whose predictions are verified by the global model are included in the set of labelled data.
4. The personalized model is retrained on the expanded set of labelled data.
5. Steps 2, 3 and 4 are repeated until all data instances are included in the set of labelled data.

If the global model exhibits poor accuracy on initially labelled data, then the inclusion of data instances in the set of labelled points will be based on confidence scores provided by the personalized model. In the case of ensemble-based models, the member of the ensemble having the highest accuracy on initially labelled data will be selected to verify predictions made by the personalized model.

#### *3.2.2.4 Advancement 4 - Unsupervised cancer-care exploratory data analytics and outlier detection in federated learning settings*

ASCAPE will also support federated learning of unsupervised models in order to enable federated approaches to data analytics and outlier detection (detection of data instances strongly deviating either from the rest of local data or unknown data instances that were previously used to train the global model). Cancer-care exploratory data analytics in ASCAPE will be enabled by clustering, association inference and outlier detection algorithms adapted for incremental and semi-concurrent federated learning schemes. Outliers detected in local data (data residing on ASCAPE edge nodes) will be excluded when training personalized models and updating global supervised and semi-supervised models. ASCAPE outlier detection methods will be based on federated models learning latent lower-dimensional data representations (e.g. deep autoencoders). Unsupervised cancer-care data analytics techniques based will also enable measuring similarity between data instances coming from different data federation partners without any data exchange. Consequently, ASCAPE will be able to identify clinics having patients with similar characteristics and to self-tune its ensemble-based predictions and model rollback options.

#### *3.2.2.5 Advancement 5 - Extendable base of federated models for cancer-care predictions*

The ASCAPE platform will enable data federation partners not only to use and update existing predictive models available in the ASCAPE cloud. They will be also able to register new models initially trained on their own local data. This will be achieved by keeping a database of federated models in the ASCAPE cloud. For each model, the database will contain its specification in terms of features and their types, the current state of the model and its history (i.e., the state of the model after each update with

characteristics of updates). The initial models will be trained using continuously delivered datasets provided by clinical partners in the ASCAPE project. Thus, the cancer-care knowledge contained in the ASCAPE platform may continuously grow in two directions - by updating existing predictive models (knowledge growth in terms of accumulated medical evidence for a particular QoL indicator or intervention enabling more accurate predictive models) and by creating new models (knowledge growth in terms of broader coverage of QoL indicators and/or interventions).

## 3.3 Homomorphic encryption (HE) for healthcare

### 3.3.1 Current State of the Art

**Privacy-preserving techniques for machine learning.**
Various privacy-preserving techniques have been developed in the past few years, addressing the balance between data utility and privacy [65]. Several approaches have been introduced and applied in machine learning based applications, for example SMPC (Secure Multi-Party Computation), DP (Differential Privacy) and HE (Homomorphic Encryption). While ensuring that data privacy is maintained, these techniques don't hinder the use of machine learning based methods for data analysis and prediction. Overall, promising results have been obtained by employing these techniques, but adoption and applicability typically depend on the use case. Since each technique has certain vulnerabilities and strengths, a trade-off is performed between performance and privacy, or between utility and privacy.
Researches have tried to address the issue of preserving data privacy in ML-based analysis relying on HE data. By definition, HE data can still be manipulated while being encrypted. Thus, data privacy is ensured while a third party processes the information in its encrypted format, without being able to understand it. With this approach, full utility of the data can be maintained, since the mathematical structures underlying the data are preserved. In the context of HE, data is collected in a centralized location, removing the communication bottleneck encountered in case of SMPC. The initial developments of HE [66] lead to a large computational overhead, making their use in ML applications impossible. Recent developments have led though to a number of ML solutions which are privacy-preserving [67], [68], [69]. One method employs a cryptosystem based on HE, which allows for standard operations to be performed on ciphertext data [69]. In case of computations that cannot be performed on the encrypted data, a communication between the server and the data owner was necessary. In a different approach the interaction between the individual parties is fully removed by employing polynomial non-linear functions with a low degree [67]. Therein, a neural network trained apriori is applied on data that is encrypted to obtain an encrypted result. YASHE [70] is employed as encryption scheme, i.e. operations on floating-point number cannot be performed (a conversion from floating-point to integer numbers is required). The usability of this approach is further limited by the fact that the computational complexity is large in case of complex networks. CryptoDL [68] attempted to mitigate this issue by employing low-degree polynomials for the approximation of functions that are non-linear. We note that privacy preservation during model training is not addressed in these schemes. Moreover, the most important disadvantage of the above-mentioned approaches is the computational

complexity: as the networks become deeper, significantly longer runtimes are obtained. Furthermore, if an approximation is performed for addressing neural network model non-linearity, performance does not necessarily improve. The majority of HE based approaches do not ensure the best prediction performance, since activation functions employing polynomial approximations are used.
To still obtain high levels of privacy and accuracy, researchers started to combine different approaches. For example, DP and SMPC techniques were combined for performing model training in a collaborative and privacy-preserving way [71]. Therein SMPC is employed to enable ML-based analysis when multiple parties hold the data, while DP addresses data security. Experiments revealed that performance decreases when large networks were trained. Alternatively, SMPC was employed for non-linear functions within a solution relying on neural networks operating on HE data [72]. Finally, in a different approach the linear operations are based on HE, while SMPC is employed for computing the activation functions [73].

**Homomorphic encryption.** Since Gentry first introduced the FHE (Fully Homomorphic Encryption) scheme [66], a large number of variants derived from the initial strategy have been proposed [74]. The majority of the schemes provide high security, but also a significant computational overhead, while only a small number of operations are allowed to still be able to perform the decryption. Thus, real-world usage is limited. Two major challenges arise for the use of deep learning for data analysis: when compared to the plaintext versions computations are orders of magnitude slower, and noise accumulates with each operation. Furthermore, no scheme is available for operating directly on floating point numbers.
Various HE libraries, open-sourced, have been developed [75]. SEAL (Simple Encrypted Arithmetic Library), developed by Microsoft offers support for the Cheon-Kim-Kim-Song (CKKS) [76] and Brakerski/Fan-Vercauteren (BFV) schemes [77]. HELib, developed by IBM [78], is based on the Brakerski- Gentry-Vaikuntanathan (BGV) scheme [79]. HELib does not support operations on floating-point data. SEAL addresses this limitation by exploiting the properties of the CKKS scheme and specifically its capability to rescale numbers without affecting the plaintext values. Data is represented by polynomials with coefficients that are integers, and a parameter scales the floating-point parameters, affecting computational accuracy. Noise is introduced in both SEAL and HELib, leading to a limited number of ciphertext operations that can be performed. This has led to the development of techniques for managing the noise, i.e. to keep the noise value lower than a given threshold, avoiding ciphertext corruption. Bootstrapping, a computationally expensive procedure is used by HELib to allow for an unlimited number of operations. Within SEAL, the number of operations in the computations has to be estimated, and then, an error reduction technique that is scale-invariant is employed. For both SEAL and HELib only multiplication and summation can be performed fully homomorphically. Non-linear operations need to be substituted by polynomials with a low degree and divisions are not supported. Due to these disadvantages, the topology of neural networks employing such encryption schemes, is significantly constrained, leading to lower prediction accuracy in privacy-preserving deep neural networks [80].
Different approaches have been introduced in the past, based on PHE (Partially Homomorphic Encryption). A solution based on PHE typically specializes on certain

operations, required for a specific use case. Significant advantages are obtained in terms of execution time [81]. Herein, we refer to the Paillier scheme [82], where a multiplication performed on plaintext data corresponds to an addition performed on ciphertext data, and to the ElGamal scheme [83] which natively is multiplicative, but can be transformed into an additive scheme. Other PHE schemes which can be employed in real-world applications are: deterministic encryption [82] (encrypted data equality checks), Goldwasser-Micaly [84] (XOR operation), searchable encryption [85] (keyword search), and order-preserving encryption [86] (encrypted values sorting).

A different interesting approach proposed in literature is the AHEE (Algebra Homomorphic Encryption) scheme [87]. This scheme is homomorphic wrt to algebraic multiplication and addition. In terms of computational complexity, it is similar to ElGamal and Paillier, but allows for both multiplication and addition to be performed using the same scheme. The main disadvantage is that only relatively small integer values can be encrypted; this applies also for the ElGamal and Paillier schemes. This is due to the fact that the encryption relies on exponentiation (the plaintext value represents the exponent), potentially leading to overflow even when a library allowing for multi-precision arithmetic is employed. Concretely, the largest number to be encrypted is 103 when integers represented on 1024 bits are used. Furthermore, one cannot evaluate whether an encrypted value is too large for a specific operation, and this represents a major drawback when performing operations on ciphertext data.

To be able to employ deep learning models in a privacy-preserving fashion, the cryptographic scheme must allow for computations to be executed on floating point numbers. The standard approach is to employ an encryption system which encodes floating point numbers as a series of integers [88]. Such an approach is of limited use in real world applications, since even some basic operations are difficult to perform on encoded data. Moreover, not only the utility of the data is affected, but also the accuracy of the results.

Hence, we can conclude that several HE schemes have been introduced, which meet security specifications. However, the majority of these methods cannot be used in real-world applications since execution times decrease by orders of magnitude compared to the computations performed on plaintext data. As a result, simpler encryption schemes relying on linear transformations have been introduced. Although offering weaker security [89], currently these encryption schemes are the only ones that offer a certain level of privacy-preservation in real-world scenarios.

The methodology to be further developed (see next sub-section) is formulated from a version of the matrix-based HE scheme proposed previously [90]. Compared to other HE schemes exploited in solutions allowing for privacy-preservation in deep neural networks [84], [85], [88], the MORE scheme is non-deterministic (when the same plaintext value is encrypted using a certain key different ciphertext values are obtained) and noise-free. Hence, one can perform an infinite number of operations without losing accuracy. Furthermore, all four elementary arithmetic operations can be performed on ciphertext data.

### 3.3.2 ASCAPE Beyond State of the Art

Motivated by the need for rational number arithmetic in homomorphic encryption and given that the MORE encryption scheme cannot be used directly on rational numbers due to its weaker security, in ASCAPE we propose an improved encryption scheme.

#### *3.3.2.1 Advancement 1 – Hybrid MORE encryption scheme*

The proposed method is based on polynomial evaluation maps, i.e., a rational number is first represented as a polynomial and then the coefficients are encrypted using the MORE scheme. An advantage of this approach is that the resulting polynomial can be forced to have integer coefficients, therefore enabling the possibility of using the standard MORE or even classic homomorphic encryption schemes.
The proposed encryption algorithm encodes a floating point message $m$ into a plaintext polynomial as follows: $P(x) = a_n x^n + \ldots + a_0 x^0 = m$, where $x$ is a secret random number.
Following the MORE strategy, each of the polynomial coefficients of $P$ is encoded into a ciphertext: $C(a_i) = SA_iS^{-1}$, where $S$ is the secret key and $A_i$ the $2x2$ constructed matrix from the polynomial coefficient $a_i$ and two random parameters $r_1$ and $r_2$ placed on the main diagonal and off-diagonal respectively. The evaluation map $e_k$ is a function from $R[x]$ to $R$. For any polynomial $f \in R[x]$ and $k \in R$, we set $e_k(f) = f(k)$. This is a ring homomorphism.
Let $f(x) = a_n x^n + \ldots + a_0 x^0$, and $f(x) = b_n x^n + \ldots + b_0 x^0$, where $a_i, b_i \in R$, we have:

$$\begin{aligned} e_k(f+g) &= e_k\big((a_n + b_n)x^n + \cdots + (a_0 + b_0)x^0\big) \\ &= (a_n + b_n)k^n + \cdots + (a_0 + b_0)k^0 \\ &= a_n k^n + \cdots + a_0 k^0 + b_n k^n + \cdots + b_0 k^0 \end{aligned}$$

Hence, $e_k$ is an additive group homomorphism and same holds true for multiplication and division.
To recover the floating point data that represents the initial plaintext message $m$, one needs first the secret key $S$ to decrypt each coefficient of the polynomial but also the secret number $x$ on which to evaluate the decrypted polynomial.
Therefore, the proposed Hybrid MORE scheme is fully homomorphic with respect to algebraic operations: addition, subtraction and multiplication. Performing division, however, is slightly more complicated as it requires polynomial division which typically results in a quotient $Q$ and a reminder $R$. A simple solution consists in representing a cyphertext as a fraction of polynomials $\frac{A(x)}{B(x)}$ rather than a single polynomial. Division can then simply be performed by multiplying with the inverse fraction. The drawback is that the addition of two ciphertexts will require a scaling operation for the fractions, to enforce the same denominator.

#### *3.3.2.2 Advancement 2 – Handling non-linear functions / limiting polynomial growth in the Hybrid MORE encryption scheme*

Other challenges arise when non-linear functions are applied on ciphertext data. Constructing a polynomial $P$ requires a secret random number $x \in X$ such that $P(x) = m$. A possible solution for enabling non-linear functions over ciphertext data can be

formulated as follows: knowing the domain $X$ of the secret number, one can sample $N$ numbers from domain $X$, apply the non-linear function, and use a simple regression task to fit the new data, and, hence, provide the new encrypted polynomial.
Note that, as for any polynomial based solution, to be able to perform an unlimited number of operations, a mechanism to observe and limit the growth, becomes mandatory. Moreover, by integrating the polynomial representation in the scheme, one could force the coefficients to be in the $Z$ domain, which implies the encryption of integer values, and, hence, the original scheme as proposed by Kipnis et al. [90] could be further employed. However, such an approach will introduce some noise due to the approximation mechanism.
Within ASCAPE, the newly developed hybrid MORE encryption scheme will be employed together with AI algorithms [65] to enable AI based privacy preserving decision support.

## 3.4 Privacy-aware AI for healthcare based on epsilon-differential privacy

### 3.4.1 Current State of the Art

Privacy is one of the core concepts in all data intensive applications which include collecting different kinds of data from data subjects (which are common human individuals). Consequently, many privacy preserving techniques were developed in order to protect the private information of involved subjects. The term Personally Identifiable Information (PII) [91] is used for any data that can be used to identify individuals, either directly or by combining the information with other data. The most basic process of protecting privacy is the process of data anonymization which includes either encrypting or removing PII from a dataset in order to protect the privacy of the individuals.
However, it was shown [92], [93] that the identity of individuals can be quite easily detected even from the published dataset with identifying information removed. According to that research 87% of the American population can be uniquely identified by date of birth, gender and postal code. Accordingly, several better solutions were developed for privacy preservation:

- *K*-anonymity: all combinations of quasi-identifiers must be repeated at least *K* times in the database [94]. Quasi-identifiers are groups of attributes that can be used jointly to identify a person or a group of persons (date of birth, gender and postal code, from the previous example). In this case, the probability of identification of a particular person is *1/K*.
- *L*-diversity: for each group of individuals with the same quasi-identifiers there must be at least *L* different values for each confidential attribute [94]. This approach introduces diversity in the data sharing some common attributes.
- *T*-closeness: the distance between the distribution of the confidential attribute in the group of the persons sharing same quasi-identifiers and the distribution of the attribute in the whole data set is no more than a threshold *T* [94]. This principle states that the distribution of sensitive attributes in some groups of individuals must be approximately same as the distribution of that attribute in the whole dataset.

All mentioned exact principles have some disadvantages. Firstly, the structure of the dataset must be significantly changed to satisfy these requirements, which is commonly not feasible. Secondly, in order to adequately define quasi-identifiers data owners/curators should know the level of knowledge of possible attackers. If that knowledge is unknown (which is usually the case) the protection can fail completely.
One of the possibilities to cope with the above mentioned problems is the utilization of differential privacy (DP) mechanism [95], [96], [97]. It represents a powerful privacy protection technique which does not require any insights into the structure of knowledge of attackers, nor does it require the reorganisation or restructuring of the dataset. Differential privacy is based on the idea that the outcome of the query posed to protected database is essentially equally likely independent of whether any individual joins or refrains from joining the database [98]. In such a way the private data about particular participant is absolutely protected since the system returns the same result (with the same probability) weather that participant was involved in analysis or not. Differential privacy mechanism tries to minimize knowledge about individual, while maximizing the knowledge about whole population.
Differential privacy is not a unique algorithm, but a methodology which can be used for developing a plethora of algorithms [99]. More formally, a randomized algorithm $A$ is ε-differentially private if for all neighbouring databases $D_1$ and $D_2$ (databases which differ only in one row) and for all sets Ω of possible outputs the following condition holds [95]:

$$Pr[A(D_1) \in \Omega] \leq \exp(\varepsilon) \cdot Pr[A(D_2) \in \Omega]$$

Simplified, for a small value of $\varepsilon$, this equation shows that the probability of the output Ω of the algorithm $A$ with database $D_1$ is nearly the same as the probability of the same output with the database $D_2$. Databases $D_1$ and $D_2$ differ only in one row, so for a small value of ε the adversary cannot learn anything about an individual record regardless of whether the record is present or absent in the analysis.
Parameter $\varepsilon$ is called the privacy budget [100]. This parameter controls the level of privacy of the algorithm $A$. Smaller values of $\varepsilon$ mean stronger privacy. The value 0 represents total privacy, but the usability of such algorithm is minimal since in that case algorithm represents pure randomness.
There are three main mechanisms that can be used for the implementation of differential privacy: Laplace mechanism, Gaussian mechanism and Exponential mechanism. Laplace mechanism is most commonly used for numeric types of queries. The idea behind Laplace mechanism is simple: it consists of adding Laplacian noise to the answer of query q before returning the result to the data analyst. The Laplacian noise follows Laplace distribution. The scale (parameter b in the distribution) is set to sensibility of query divided by $\varepsilon$. With such set-up Laplace mechanism satisfies ε-differential privacy condition.
The goal of ML is in a way similar to the goal of privacy: the learner wants to discover some rule that explains the whole dataset. Ideally, this rule should be applicable not only on existing data, but also on some future data from the same domain. Therefore, ML tries to capture distributional information about the data set in a way that does not depend on any single data record. That is exactly the goal of private data analysis. Unfortunately, combining differential privacy with data mining and machine learning algorithms is not an easy task. The main reason is probably the fact that it is possible to reveal the private information even from the output of the ML models [101]. This

kind of attack is called model inversion attacks. To overcome this problem the Privacy INtegrated Queries platform (PINQ) is proposed [102], where the ML model is trained with already privately protected data. In case of differential privacy this means with the data with already added noise [103]. The main challenge with the proposed framework is to find an appropriate trade-off between the privacy and the utility of the ML algorithm.
Differential privacy has found its application in numerous fields, for example:

- The US Census Bureau implemented DP in their OnTheMap project to ensure privacy for population data. OnTheMap application gives researchers access to agency data.
- Apple used DP for three different purposes: a) for discovering popular emojis, b) for identifying resource-intensive websites in Safari, c) for discovering the use of new words.
- Microsoft used DP to hide the true location of individuals in their geolocation databases.
- Uber uses DP in their data analysis pipeline and other development workflows.
- Google uses DP in their RAPPOR project, which is used to report usage statistics for Google Chrome.

Although DP has been applied in various, very important fields and applications, it is still rarely applied to medical data [104]. Fortunately, some recent researches [105] try to change that.

### 3.4.2 ASCAPE Beyond State of the Art

The guarantees of differential privacy are rather strong but can come at the expense of accuracy [102] especially when dealing with ML models. In such cases there is always a trade-off between data privacy, accuracy of ML models, and the dataset size [103].

#### *3.4.2.1 Advancement 1 – Analysis and tracking of privacy tuning parameter*

ASCAPE will develop a framework for application of differential privacy AI and machine learning method pools for analysis of medical data. The specific data structures will be considered in order to develop optimal differential privacy architectures, i.e., optimal values of noise-adding sub-blocks within the algorithms. The privacy-accuracy trade-offs which arise in medical applications will be carefully analysed in ASCAPE framework and tracked in order to offer to the patients and hospitals appropriate service and the insights into tuning the privacy epsilon-factor for the targeted application.
The AI model training with DP will be used on edge-nodes of the healthcare providers having sufficient resources on the local edge nodes, for instance as part of a federated averaging approach to federated learning. For healthcare providers with limited processing resources, i.e. in clinics with not sufficiently powerful edge nodes to perform local model training, DP shall be used before the central cloud model training for instance as part of a federated stochastic gradient descent based federated

learning. The models obtained from DP-based federated learning can be used by all pilot sites.

#### *3.4.2.2 Advancement 2 – Privacy tuning parameter recommendation*

The ASCAPE system will contain a tuning parameter (ε) which will control the trade-off between the amount of information leakage on the one hand and the quality of the produced analytics on the other hand. The value of tuning parameter value can be controlled by the platform user/data provider (e.g., an individual patient, a hospital, a physician, etc.) based on their decided level of information leakage. However, it's hard to expect that the average user has the required knowledge about DP mechanisms to adequately select the value of ε. Therefore, the ASCAPE system will propose the optimal value of tuning parameter to the end users based on the quality and quantity of their local data. Surely, the proposed value could be changed by the end users, but the idea of this proposition is to relieve the patients and physicians of additional efforts during usage of the system.

# 4 Requirements Specification Methodology

The approach taken in specifying the requirements was affected by a number of factors each with its own import in the process: scope, constraints, and purpose.

## 4.1 Scope

The scope of the present deliverable is to be understood in the context of the overall ASCAPE workplan.
In WP1 here are three tasks and two deliverables that touch on the subject of specifications and requirements.

- The present deliverable, D1.1, due in month 6 of the project, stems from work on the ASCAPE framework specifications and requirements task (T1.1) which focuses on the ASCAPE framework.
- The follow-up deliverable, D1.2 “ASCAPE Data Determinants and piloting validations”, due on in month 8 of the project, focuses on requirements pertaining to the ASCAPE pilots and stems from work on two pilots-focused tasks:
  - Task 1.2. Data Determinants affecting Quality of Life for cancer patients
  - Task 1.3. In depth analysis for ASCAPE pilots and Quality of Life use-cases

The rationale of the workplan is to address aspects pertaining to the general framework in the present deliverable, D1.1, while allowing more time for work in finalising pilot specification and data determinants to be addressed in the following WP1 deliverable, D1.2.
Consequently, the present deliverable is not meant to address details of the pilot studies but focus on the overall framework. The ASCAPE framework itself is an, otherwise domain agnostic, AI framework focusing on applications in the medical applications where its technological innovations that can challenge the orthodoxy of the Iron Triangle of Health orthodoxy and build AI knowledge on the cloud based on large quantities of sensitive medical data from a large number of health providers without requiring that such data leaves the health provider’s IT infrastructure.
Accordingly, requirements are spread across D1.1 and D1.2, with D1.1 focusing on non-functional requirements and generic outlines of functional requirements and D1.2 with details about what aspects of quality of life are to be monitored for breast and prostate cancer, what will be the data collection framework in each pilot study (patient records, standardised questionnaires, mobile apps, website forms, wearables etc.) and pilot-specific use cases.
In summary, general framework requirements are within the scope of the present deliverable, whereas requirements for particular applications of the framework (e.g. the four pilots and participation in the open call) are not. This separation of concerns helps ensure the current deliverable clarifies how the ASCAPE framework can provide a foundation not only for the pilots but for a variety of medical applications, with the pilot studies acting both as a demonstration and a testbed of its capabilities.

## 4.2 Project Commitments

The ASCAPE framework could address a number of medical issues, indeed also issues unrelated to medicine, but as per the aims of the ASCAPE project, as recorded first in the Proposal and subsequently in the Grant Agreement, it will be considered as a novel trustworthy big-data AI platform for supporting cancer patients, focusing on their QoL.
Commitments made are not only thematic (focus on cancer patients' quality of life), but also technological and procedural. ASCAPE promises to deliver a platform that:

- Enables the centralised gathering of knowledge from local data using advanced methods in ML including federated learning and ML on homomorphically encrypted data while ensuring that local data are not themselves transmitted outside the confines of the health provider's IT infrastructure (offering instead the aforementioned methods for achieving the goal of building global knowledge from local data)
- Offers the ability to take advantage of that centralised knowledge in providing predictions and health intervention suggestions for patients whose data are maintained locally at their health-provider's infrastructure
- Aims to inform the medical opinion of doctors (in a similar manner that studies and guidelines do), not to provide medical advice directly to patients.

The requirements will reflect the above fundamental design principles. This includes:

- Security requirements that will re-enforce the privacy of data and trustworthiness of the ASCAPE platform
- Performance requirements that will address both the necessity for instant predictions and the necessity for the ASCAPE platform to allow training on the basis of large quantities of data
- Functional requirements that take into consideration the principle that local data are not gathered centrally; this has consequences, for example with regards to what visualisations can be provided (see Section 7.1).

## 4.3 Purpose of Framework Requirements

A final consideration in planning the content and approach towards the ASCAPE framework requirements developed herein, no less important than their scope and the relevant commitments undertaken by the Consortium in delivering ASCAPE, is an understanding of what the ASCAPE framework represents in the overall vision for ASCAPE – and equally importantly what it does not.
The ASCAPE framework is not to be seen as a product. This is because an exploitation path whereby ASCAPE would attempt to become a complete software solution, say, for cancer clinics would place it in the market as a competitor to existing solutions meaning that the benefits of the project delivered to cancer patients would be limited by the piece of the market pie ASCAPE would be able to claim.
Instead, the exploitation path chosen is one that turns the ASCAPE outputs into valuable resources every software vendor with an existing software solution would be able to take advantage of and the ASCAPE partners into potential strategic allies. In what is envisaged to be the typical ASCAPE exploitation scenario, the aim will be for

software vendors whose software is used by a number of Medical Care Providers to create ASCAPE-powered versions of their existing software; this way the benefits of ASCAPE can both be delivered as updates to existing software and the software vendors' sales teams efforts will, as a side-effect of promoting the newer version of their software to existing and/or new customers, help further spread of the benefits of ASCAPE.
It is envisaged that ASCAPE-powered versions of existing software solutions will be created by re-implementing and/or integrating part of the ASCAPE edge infrastructure (the part that is meant to resides on the clinic's infrastructure) using the ASCAPE cloud (in order to realise the vision of maintaining and using creating a global source & sink of AI-based knowledge)
The ASCAPE Dashboard (developed as part of the Pilots work package, WP4) will aim to showcase ASCAPE technology in the context of the ASCAPE Pilots and provide guidance to software vendors about how the AI functionality can be integrated visually and functionally into their products.
This vision for the future of ASCAPE played an important role in steering the process to determining the focus and aims of the requirements. For instance, functional requirements focus on the ASCAPE-specific functionality to be showcased in the ASCAPE Dashboard and implemented in the UIs of ASCAPE-powered software solutions, whereas access control, user management, data entry, reporting and other non-ASCAPE-specific functionality are not covered, both because ASCAPE does not contribute anything of interest in these areas and because existing software solutions that we would like to see ASCAPE-powered versions of already implement their own functionality in these areas.

## 4.4 Requirements Specification Process

The requirements specification process was focused, with known targets, known constraints and a well understood scope. Setting the aforementioned parameters was non-trivial, but consensus was built gradually between technical and pilot partners in the course of reaching a common understanding of the project's objectives and the ways its benefits could be demonstrated in the pilot studies.
The ASCAPE framework requirements specification process unfolded in two phases:
**Phase I (mid M2 to mid M3): Initial System Requirements Gathering**
The first phase of the requirements specification process was initiated at the Kick-Off meeting in late January 2020. It covered both general ASCAPE framework requirements (T1.1/D1.1) and pilot-specific requirements (T1.2&T1.3/D1.2) concurrently, with most progress, inevitably, made on the former (general ASCAPE framework requirements) as the later (ASCAPE pilot-specific requirements) required significantly more deliberation—exactly as the project workplan had foreseen.
Initial requirements gathering took place over the course of approximately two months of intensive partner-internal deliberations, bilateral or multi-lateral communication between the WP leader, the Scientific Coordinator and partners needed to discuss specific topics, as well 6 online hour-long weekly discussions (from 14 February to 20 March 2020 inclusive) between the technical and pilot partners (and the legal partner on two instances).
Key achievements and outcomes included:

- An understanding of the commitments the ASCAPE Consortium has undertaken
  - with respect of the aims to be achieved
  - with respect to the technologies to be used / further developed
- An understanding of:
  - the reliance of the success of the pilot studies (and, by extension, the project as a whole) on the data that will be made available to the ASCAPE AI
  - the kinds and amounts of data that will be made available
  - the level of AI-based functionality that can be made available at different times in the project given the expected availability of data to enable said functionality (e.g. predictions and intervention suggestions with sufficient confidence levels)
- An understanding of the scope of D1.1 (and of D1.2)
- An understanding of how different exploitation options affect the project's ability to achieve its aim of democratising knowledge derived from sensitive data and how the choice of exploitation affects what the ASCAPE framework requirements process should put emphasis on
- An initial understanding of who the users of ASCAPE will be during the pilot phase of the project and in subsequent ASCAPE results exploitation:
  - Doctors (Not Necessarily Cancer Experts)
  - Patients (Indirectly: they provide data but do not directly receive predictions or intervention suggestions by ASCAPE as ASCAPE does not replace but informs their doctors' medical opinions)
  - IT System Administrators
- An initial understanding of
  - Basic functional requirements among all partners and
  - Basic security requirements (including requirements on the treatment and of sensitive)
  - Basic performance requirements (and the separation possible between interactive and batch processing)

**Phase II (M4-M6): ASCAPE Framework System Requirements Specification and Refinements**

In Phase II, focus was on ASCAPE Framework Requirements with a clean separation between these (recorded herein) and Pilot-Specific Requirements (to be recorded in D1.2).

System Requirements Specification begun with decisions about how the requirements are to be recorded and organised and decisions about whether use cases and scenarios are to be used. Based on preparatory work by the WP1 and T1.1 leaders and over the course of two weeks (23 March to 3 April 2020) the following decisions were taken:

- It was decided that use cases will be used as the basis for functional and, where appropriate, non-functional requirements and that these use cases would be generic, addressing the functionality provided by the framework abstracting away from the details of any particular application of the framework (including the four ASCAPE pilots).

- It was decided to use structured natural language for recording the requirements, using the templates of Section 4.5.
- The overall structure of use cases, functional and non-functional requirements was, provisionally, finalised.
- It was also decided to have use scenarios, albeit not with the express purpose of being part of the requirements specification, but rather as a means of introducing the aims and vision of the ASCAPE.

Even though the ASCAPE framework does not interact directly with doctors, it needs to be capable of providing functionality to ASCAPE-powered systems which doctors will be able to use. Within this project, the ASCAPE Dashboard for Doctors will be created to demonstrate the project's relevant AI-based capabilities and various ways of utilising the ASCAPE infrastructure in conjunction with pilot partner's existing IT systems, in effect leading to the creation of the first ASCAPE-powered systems. The main users of such systems will be doctors looking after cancer patients. The project's rationale, supported also by the findings of the requirements gathering process is that ASCAPE should assist doctors in their goal to better look after their patients' health (focusing on quality of life aspects), not try to replace them, nor to bypass them.
Therefore, ASCAPE AI is meant to provide predictions and intervention suggestions to doctors and it will be up to the doctor to use this information as they see fit, using their own medical expertise and knowledge of the patient's case. Yet unlike clinical studies, guidelines and other sources doctors have been using to supplement and shape their knowledge, ASCAPE will be able to offer patient-specific results. The requirements gathering process confirmed the Consortium's understanding of doctor's needs, namely the need to have quick access to all relevant information, which resulted in a functional requirements specification focused on providing doctors with the information they need to focus on as well as very efficient means of exploring alternative intervention recommendations:

- When a Doctor visits the Patient's record, the most important information about that patient should be presented to the Doctor, prioritised by importance and with minimal, if any, additional input; additional user actions will be required for exploring alternatives and/or obtaining explanatory details about a result provided by ASCAPE's AI
- When the Doctor enters the System, ASCAPE is to provide a list of patients the cases of which ASCAPE AI believes the Doctor should review, either because there is an actual or predicted deterioration in their health and quality of life or because ASCAPE predicts an intervention will have a significant effect

Patients were not involved in the requirement gathering, as patient recruitment will take place further down the project. For the purposes of the present deliverable, they are indirect users of the ASCAPE platform (as the ASCAPE result are presented to their doctors) and the various means of collecting data from the patients are outside the scope of the ASCAPE AI infrastructure and the present deliverable. The exact requirements on questionnaires, mobile apps, wearables etc. are for healthcare providers to decide, as reflected also in the follow-up deliverable covering pilot specifications; what is required of the ASCAPE framework, and therefore, within the scope of the current deliverable, is that the ASCAPE framework can support interoperability with a healthcare system providing a number of data collection modalities and the use of the collected data both for the benefit of the specific patients

whence they originate and for the purpose of training ASCAPE models for the benefit of all patients for whom ASCAPE will make predictions and intervention suggestions in the future. The relevant data collection modalities were specified by clinicians in the project.
The requirements gathering process focused on collecting input from two of the pilot partners, whereas the other two pilot partners were asked to evaluate the resulting requirements. The first version of the ASCAPE framework system requirements specification was created over the course of approximately two weeks (6 – 17 April 2020) of collaborative effort on the basis of the Phase I results and the Phase II preparatory work mentioned above. The results of this work were presented and discussed in the First ASCAPE Plenary Meeting on 4 May and on a dedicated discussion on 8 May where all partners discussed the first version of the ASCAPE framework system requirements specification after having reviewed it in its entirety. This was an important milestone as it completed the move from the Phase I efforts to achieve a shared understanding of ASCAPE framework requirements to a concrete ASCAPE Framework Systems Specification (Section 6) which could be evaluated (across number of criteria, such as completeness, lack of ambiguity, usefulness, achievability, and correctness of prioritisation) and refined where appropriate. Subsequent versions of the ASCAPE framework system requirements were created in the following weeks, as part of the corresponding versions of the present deliverable.

## 4.5 Use cases and requirements specification templates

We have defined templates for use case scenarios and requirements in order to have a harmonised structure and enforce a standard layout and look across all the collected scenarios and requirements. The templates provide the framework that brings together common elements, gives a unique reference ID to every scenario and requirement to facilitate the linkage between them and encourages repeatability and efficiency.
The above templates for requirements and scenarios give textual descriptions inspired by a standardized formal language in RFC2119 [106] to describe among others: path of events, trigger, pre-conditions and post-conditions, rationale, and keywords. We often highlight in block letters MUST, SHOULD and COULD/MAY. These should however not be confused with the similar keywords that we use for the priority of accomplishment of a requirement, i.e., “Must have”, “Should have”, “Could have”.

### 4.5.1 The use case scenarios template

The defined use case scenario template is as follows:

| | |
|---|---|
| ID | *A unique id distinguishing this use case from any other. To form use case IDs the following scheme should be used:*<br>*<Prefix>.<number>*<br>*where*<br>*<Prefix> := HP \| PT \| ADM* |

| | |
|---|---|
| | *The use of the one of the above prefixes should indicate the primary actor of the use case. The intended meaning of the individual values is as follows:*<br>• *HP: Healthcare Providers*<br>• *PT: Patients*<br>• *ADM: Administrators* |
| Name | *A short string indicating the meaning of the use case.* |
| Description | *A brief summary outlining the overall purpose of the use case and the interaction taking place between the ASCAPE platform and the use case actor(s).* |
| Actors | *The stakeholders who will interact with the ASCAPE platform in the context of the use case. Only actors that have a DIRECT interaction with the ASCAPE platform as part of the specific use case should be listed here.* |
| Preconditions | *All conditions that must be satisfied prior the commencement of the interaction described by the use case.* |
| Trigger | *The event/circumstance(s) that will trigger the interaction described by the use case.* |
| Main path | *The typical path of steps that should be taken to realise the interaction between the ASCAPE platform and the use case actors that this use case describes. The steps should be listed in the exact sequence in which they occur and be numbered in a way that indicates this sequence (i.e., 1. (for first step), 2. (for second step) etc.) Steps should be atomic (i.e., a step should be a single action that is taken by either the external actor or the system as part of the interaction) and indicate with clarity who is responsible for taking the step (i.e., the system or an external actor).*<br>*A step may involve the extension of the use case by another use case or the inclusion of another use case within this use case. If this is the case, the step will be an extension or an inclusion point, respectively.*<br>*A use case A needs to be extended by another use case B, if the interaction described by B should take place in the context of A under certain conditions. These conditions must be clearly described.*<br>*A use case A needs to include another use case B, if the interaction described by B should take place in the context of A in all circumstances.*<br>*An extension point will be introduced by the special keyword: EXTENDED BY <use-case-id> UNDER CONDITION <condition>.*<br>*An inclusion point will be introduced by the special keyword: INCLUDES <use-case-id>* |
| Alternate path | *Alternatives to the typical steps taken to realise the interaction described by the use case. A use case may have one more alternate paths.* |

| | |
|---|---|
| | *Each alternate path may involve one or more alternative steps all of which must be associated with steps in the Main Path.*<br>*Alternate paths and their steps must be numbered according to the following scheme:*<br>*AP<number>. AP_STEP<number>.MP.<number>*<br>*where*<br>• *AP<number> is the unique identifier of the alternate*<br>• *AP_STEP<number> is the unique identifier of the individual step within the alternate*<br>• *MP.<number> is the identifier of the step in the main path that will be substituted for by the alternate path step* |
| Postconditions | *All conditions that must be satisfied upon the completion of the interaction described by the use case.* |

### 4.5.2 The requirements template

The requirements template is as follows:

| | |
|---|---|
| ID | *A unique ID for this requirement/assumption* |
| Name | *A title/short name for this requirement/assumption* |
| Priority of accomplishment | *One of the following:* ***Must have****: The system must implement this requirement to be accepted.* ***Should have****: The system should implement this requirement: some deviation from the requirement as stated may be acceptable.* ***Could have****: The system should implement this requirement but may be accepted without it.* |
| Description | *Specify the intention of the requirement/assumption* |
| Rationale | *If the description is not descriptive enough, this entry gives a justification of the requirement/assumption. Otherwise this entry will be filled with N/A.* |
| Supporting materials | *If applicable, give a pointer to documents that illustrate and explain this requirement/assumption. Otherwise this entry will be filled with N/A.* |

# 5 Use Cases

The use cases in 5.1 and 5.2 illustrate the general idea of potential interactions between the ASCAPE and healthcare providers or patients in clinical practice. There are, however, differences in how the healthcare systems are organized among different countries, differences that are also reflected on the healthcare pathways for cancer patients in each system. These differences are mostly on the way each healthcare system organizes the follow-up and rehabilitation plan for cancer patients rather than the treatment strategies themselves.
The ASCAPE is planned to be integrated into clinical practice by taking into account the different healthcare pathways for cancer patients among different countries.
As already specified in Section 4.4, the below end-users will be involved in the elicitation of the ASCAPE requirements:

I. Healthcare providers (doctors)
II. Patients
III. System Administrators

## 5.1 Healthcare providers

Use Case HP.1 – Patient Visit

| ID | HP.1 |
|---|---|
| Name | Patient Visit |
| Description | This use case focuses on the interaction of the Doctor with the ASCAPE-powered Healthcare Provider Information System in the context of a Patient Visit. |
| Actors | The Doctor, The Patient |
| Preconditions | The Patient has entered the optional ASCAPE personal data collection scheme offered by their healthcare provider (See Use Case PT.1 – Patient interaction with ASCAPE) |
| Trigger | Patient visit to the Doctor |
| Main path | 1. The Patient enters the Doctor's Office<br>2. The Doctor finds the Patient's record in the ASCAPE-powered Healthcare Provider Information System.<br>3. The System logs the fact that the Doctor has visited the Patient's record<br>4. The Patient's record, thanks to ASCAPE-integration, offers features such as:<br>a. A list of ASCAPE signals regarding the patient's Quality of Life metrics (current and/or predicted)<br>b. What-if Graphs of QoL metrics visualising the historic and/or predicted values of QoL metrics and the potential effect of proposed interventions; the default configuration when the Doctor opens the Patient's record presents a comparison of any ASCAPE-proposed interventions with the no-intervention case |

| | |
|---|---|
| | 5. The Doctor quickly peruses the QoL signals, by clicking on each one of them. There are two kinds of signals:<br>a. Purely informative signals highlight a concern about the patient's quality of life either on the basis of current data or on the basis of a predicted trajectory of one or more indicators, where the prediction can be made high sufficiently high accuracy; upon the Doctor clicking on one of those signals the graphs for the relevant metrics appear, clearly distinguishing between observations up to that day and predicted values, as well as providing confidence levels for the latter<br>b. QoL intervention suggestion signals, which appear when there is sufficiently high confidence that an intervention will improve the Patient's QoL, show information about why the specific intervention suggestion is made.<br>6. The Doctor, curious to see why another possible intervention was not recommended, uses the What-If Graphs manually to see ASCAPE's view on that intervention and then proceeds to obtain an explanation from ASCAPE about the basis on which it proposed the interventions that it did<br>7. The Doctor consults with the Patient and determines if there are any medical reasons, possibly not recorded in the patient's history, making any of the considered interventions inappropriate, updates the patient's history if necessary, and judges what the optimal course of action is appropriate for the Patient.<br>8. The Doctor discusses the recommended course of action and any alternatives with the Patient (including the wait-and-see/no-intervention option where appropriate) and shares ASCAPE predictions about each of option if they find them to be in accordance with their own medical opinion<br>9. The Doctor and the Patient agree on a course of action (which the Doctor records in the ASCAPE-powered IT system) and discuss what the Patient should do, what symptoms to look out and report to the Doctor (or a colleague of a particular specialisation where appropriate). and what expectations they should have<br>10. The visit and the agreed course of action are recorded by the System<br>11. The Doctor and the Patient renew their appointment either for after the regular monitoring interval or earlier if appropriate. |
| Alternate path | |
| Postconditions | The visit, the Doctor's access to the Patient's record and the agreed course of action for the Patient are recorded. |

Use Case HP.2 – Doctor Alert

| | |
|---|---|
| ID | HP.2 |
| Name | Doctor Alert |
| Description | This use case focuses on the ability of an ASCAPE-powered Healthcare Provider Information System to alert the Doctor about a developing health issue with the Patient |
| Actors | The Doctor, The Patient |
| Preconditions | The Patient has entered the optional ASCAPE personal data collection scheme offered by their healthcare provider (See Use Case PT.1 – Patient interaction with ASCAPE) |
| Trigger | Input from Patient or authorised health-monitoring device |
| Main path | 1. Input from Patient or authorised health-monitoring device reaches the ASCAPE Platform, is processed and is converted into an ASCAPE signal<br>2. An informative indicator prompts the Doctor to view the new task in the ASCAPE-powered Healthcare Provider Information System. The task description explains the issue identified by ASCAPE, the timestamp and the type of data that triggered the signal (e.g. wearable, self-reporting questionnaire or other) and contains a link to the patient's record<br>3. There the Doctor finds the complete list of ASCAPE signals regarding the patient's QoL, quickly examines them (e.g., system indicates different entries) and determines that it would be best to arrange a call with the Patient.<br>4. The Doctor uses the ASCAPE-provided functions of the Healthcare Provider Information System during the call with the patient, in a manner similar to what was described in Use Case HP.1 – Patient Visit. |
| Alternate path | |
| Postconditions | |

## 5.2 Patients

Use Case PT.1 – Patient interaction with ASCAPE

| | |
|---|---|
| ID | PT.1 |
| Name | Patient interaction with ASCAPE |
| Description | This use case focuses on how a Patient interacts with an ASCAPE-powered Healthcare Provider Information System. Unlike other use cases which concentrate on a short event, this use case covers the duration of a patient's interaction with ASCAPE. (The patient's interactions with ASCAPE listed herein are indicative, are not mandated nor dictated by the ASCAPE framework and may different from Healthcare Provider to Healthcare Provider and from patient to patient) |

| | |
|---|---|
| Actors | The Patient, The Doctor |
| Preconditions | |
| Trigger | |
| Main path | 1. The Patient enters the optional ASCAPE personal data collection scheme offered by their healthcare provider and supported by the ASCAPE Cloud and the ASCAPE-powered Healthcare Provider Information System:<br>a. The Patient signs the ASCAPE Data Processing Consent Form or a form with the same legal effect provided by their healthcare provider in order to benefit from the ASCAPE services<br>b. The patient provides access to data from a wearable provided by their healthcare provider (optional, but without this the patient's ASCAPE predictions and intervention suggestions will not benefit from data from a wearable device)<br>c. The patient downloads a mobile application provided by their healthcare provider and grants it access to send data (e.g. short QoL questionnaires) to the healthcare provider's ASCAPE-powered Information System (optional, but without this the patient's ASCAPE predictions and intervention suggestions will not benefit from data from the Healthcare Provider's mobile app)<br>2. The Patient provides data to the ASCAPE-powered Healthcare Provider Information System either directly (e.g. via questionnaires they fill in) or indirectly (e.g. via the wearable device) and so does the healthcare provider (e.g. information about what QoL improvement intervention they recommended to the patient)<br>a. The wearable device provides data on the Patient's physical activity and sleep quality (alternatively it could provide additional information, such as blood oxygenation, or not be used at all)<br>b. The mobile app may for regularly filling QoL mini-surveys (alternatively it could be used for allowing the patient to find information about a particular symptom and only then ask them if they have that and to what extent, or it could be used for filling more detailed QoL questionnaires or the Healthcare Provider or the patient may decide not to use an app at all)<br>c. The Patient is also asked to fill in more detailed QoL questionnaires on tablets given to them during their visits to the Healthcare Provider's premises as well as from the comfort of their home by using the Healthcare Provider's website (a number of further alternatives could be seen in practice, including having a member of staff ask the |

| | |
|---|---|
| | questions and record the answers either on the Healthcare Provider's premises or over the phone)<br>d. The Patient's personal information, medical history, cancer treatment information, doctor visits and lab results are also entered into the ASCAPE-powered Healthcare Provider Information System<br>Note: None of these data are transmitted by the ASCAPE-powered Healthcare Provider Information System to the ASCAPE Cloud, but the knowledge they contain is made available via advanced technical means in accordance with the Healthcare Provider's policies as reflected by the relevant system settings (see use case ADM-1).<br>3. The doctor can use AI-provided functionality provided by the ASCAPE-powered Healthcare Provider Information System based on the above data to consult the patient, suggest appropriate interventions, or adjust the follow-up schedule. See use cases HP-1 and HP-2.<br>4. The Patient's monitoring ends. Patient data stored on the ASCAPE-powered Healthcare Provider Information System may be subject to (partial) deletion in accordance with the Healthcare Provider's and the exact terms of the Patient's consent form (see Step 1.a). The knowledge obtained from the Patient's data remains in the ASCAPE Cloud for the benefit of other patients.<br>Steps 2 and 3 may be repeated and as data collection and medical follow-ups are part of a continuous process. |
| Alternate path | |
| Postconditions | |

## 5.3 System Administrators

Use Case ADM.1 - Control of transmission of data to the ASCAPE Cloud

| | |
|---|---|
| ID | ADM.1 |
| Name | Control of transmission of data to the ASCAPE Cloud |
| Description | This use case focuses on how the Administrator can change the setting that controls if the ASCAPE-powered Healthcare Provider Information System contributes to the centralised ASCAPE knowledge and by which means. |
| Actors | The Administrator |
| Preconditions | The Healthcare Provider Information System has been upgraded to an ACAPE-powered version; by default, it is configured not to attempt to update global ASCAPE knowledge |
| Trigger | |
| Main path | 1. The Administrator navigates to the settings page |

| | |
|---|---|
| | 2. The Administrator enables the sharing of anonymised local data with the ASCAPE Cloud by selecting the allowed method(s) (e.g. both federated learning and homomorphic encryption or only federated learning) |
| Alternate path | |
| Postconditions | The sharing of knowledge from anonymised local data stored in the ASCAPE-powered Healthcare Provider Information System with the ASCAPE Cloud is enabled and takes place using the advanced ASCAPE technological solution(s) that ensure that this happens without the local data themselves being revealed to the ASCAPE Cloud |

# 6 System requirements

System requirements are detailed specifications describing the functions the system needs to do and are divided into:

- **Functional requirements**
  Functional requirements determine the goals that users want to reach and the tasks they intend to perform. By eliciting the functional requirements, we understand why the user performs certain activities, what are his/her constraints and preferences, and how the user would trade-off between different software capabilities. The important point to note is that WHAT is wanted is specified, and not HOW it will be delivered.
- **Non-functional requirements**
  Non-functional requirements determine the restrictions on the types of solutions that will meet the functional requirements. Specification of non-functional requirements includes performance aspects, security, privacy, and general criteria that judge the operation of the system.

System requirements will be translated into the technical specifications based on which the initial design of the system architecture will be delivered. This part of the analysis will be presented in the deliverable D1.3 “Architecture definition”. While the details of architectural are to be determined therein, it is important to establish a high-level understanding of ASCAPE architecture for the benefit of understanding the different viewpoints that come into play when discussing ASCAPE.

**ASCAPE Cloud:** The shared scalable cloud infrastructure and the corresponding ASCAPE software components providing part of ASCAPE functionality (the other part being provided by the ASCAPE Edge Nodes), including maintenance of the ASCAPE shared knowledge (the global ASCAPE Deep Learning models)

**ASCAPE Edge Node:** A combination of healthcare provider hardware inside its secure network infrastructure and ASCAPE-provided software providing part of ASCAPE functionality (the other part being provided by the ASCAPE Cloud), including maintenance of the ASCAPE local knowledge (the local ASCAPE Deep Learning models), where appropriate.

**ASCAPE-Powered HealthCare Provider IT System:** A Healthcare Provider IT System enhanced with ASCAPE functionality; such functionality may be provided in the ASCAPE-supported way, by interoperating with an ASCAPE Edge Node (Figure 2), or by alternative means (where the software vendor makes a more significant investment and assumes full responsibility) such as by source-code level integration of open source ASCAPE software components built for ASCAPE Edge Nodes or by writing their own components capable of working with ASCAPE models and the ASCAPE Cloud.

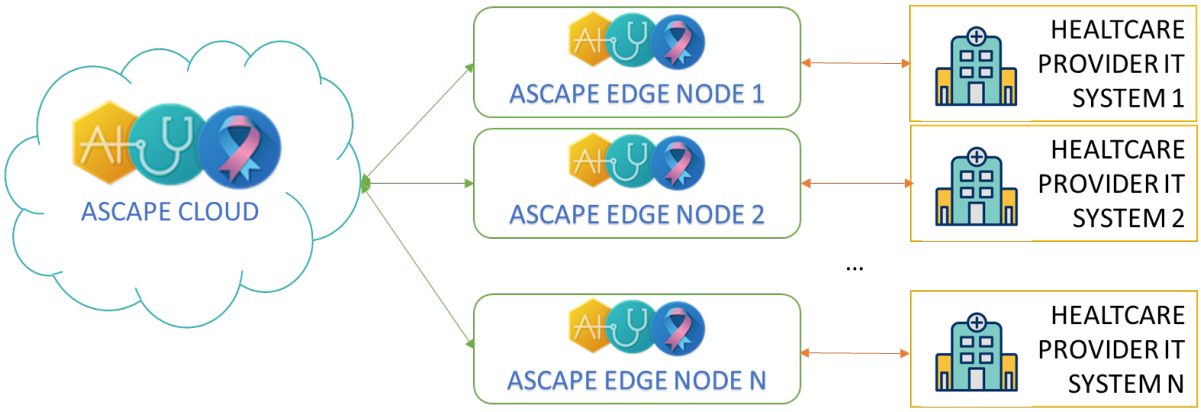


Figure 2. ASCAPE overall framework

The ultimate aim of the system requirements are to stir the development of ASCAPE in a direction that facilitates wide-scale adaptation, which can only be achieved if software vendors are convinced to provide ASCAPE-enhanced versions of their healthcare provider IT systems software. With that in mind, the system requirements, while focused on the end users, doctors, their requirements and the benefits that ASCAPE can bring to their patients, also address other issues that may affect adaptation.

## 6.1 Functional Requirements

As ASCAPE is meant as a tool for doctors the majority of functional requirements stem from the relevant use cases. As explained in Section 4, only ASCAPE-specific requirements are listed; functionality for access control, retrieval of patient records, etc. is taken for granted in the Healthcare Provider Information System. What is described below are ASCAPE-specific requirements that supplement and enhance standard functionality in such systems.

Functional Requirement FUNC01

| ID | FUNC01 |
|---|---|
| Name | Generation of ASCAPE What-If Graph for a given patient, a given metric and a (possibly empty) set of potential QoL interventions |
| Priority of accomplishment | **Must have** |
| Description | The ASCAPE-powered Healthcare Provider Information System should be able to generate ASCAPE Quality-of-Life What-If Graphs concerning a specific patient using:<br>1. its data about the patient<br>2. one or more hypotheses supplied by the Doctor for visualisation where a hypothesis will be a potential QoL-improvement intervention<br>3. the ASCAPE knowledge |
| Rationale | Supporting HP.1 & HP.2 |
| Supporting materials | N/A. |

Functional Requirement FUNC02

| ID | FUNC02 |
|---|---|
| Name | Calculation of a prominence metric for each ASCAPE What-If Graph |
| Priority of - accomplishment | **Could have** |
| Description | The ASCAPE-powered Healthcare Provider Information System could be able to calculate of a prominence metric for each ASCAPE What-If Graph for a given patient and set of potential QoL interventions determining the order in which graphs will appear to the Doctor in order to ensure the most important information is displayed first |
| Rationale | Supporting HP.1 & HP.2 |
| Supporting materials | N/A. |

Functional Requirement FUNC03

| ID | FUNC03 |
|---|---|
| Name | Generation of ASCAPE Signals for a given patient |
| Priority of accomplishment | **Should have** |
| Description | The ASCAPE-powered Healthcare Provider Information System should be able to generate ASCAPE signals concerning a specific patient as data concerning that patient reach it. Two kinds of signals should be supported:<br>1. Purely informative signals which highlight a concern about the patient's quality of life either on the basis of current data or on the basis of a predicted trajectory of one or more indicators<br>2. QoL intervention suggestion signals which are to be generated when there is sufficiently high confidence that an intervention will improve the Patient's QoL |
| Rationale | Supporting HP.1 & HP.2 |
| Supporting materials | N/A. |

Functional Requirement FUNC04

| ID | FUNC04 |
|---|---|
| Name | Instant alert upon the generation of ASCAPE Signals for a given patient |
| Priority of accomplishment | **Could have** |
| Description | The ASCAPE-powered Healthcare Provider Information System could implement functionality to email doctors or |

| | |
|---|---|
| | add a task in a doctor's to-do list, if such functionality is supported.<br><br>Depending on delivery method's level of trust either full details or merely a link to the healthcare provider's website and a patient number or signal number will be displayed. In case of a trusted medium, the body of the message should identify the patient it is about and the issue identified by ASCAPE, and should also contain the timestamp and the type of data that triggered the signal (e.g. wearable, self-reporting questionnaire or other), in addition to a link to the patient's record. |
| Rationale | Supporting HP.2 |
| Supporting materials | N/A. |

Functional Requirement FUNC05

| | |
|---|---|
| ID | FUNC05 |
| Name | All predictions of the models should be explainable |
| Priority of accomplishment | **Must have** |
| Description | In order to give insight to the decision process of a model providing a prediction, explainability techniques must be applied. The explainability tools have to suit the model architectures that are used so that the explanations are consistent with the model's internal decision process. |
| Rationale | Supporting HP.1, HP.2, PT.1 |
| Supporting materials | N/A. |

Functional Requirement FUNC06

| | |
|---|---|
| ID | FUNC06 |
| Name | The medical intervention suggestions of the models should be explainable |
| Priority of accomplishment | **Must have** |
| Description | The explanations of the explainability models should summarize the decision process in a way that is easily and immediately understandable for a person without a medical background or knowledge about the model. This ensures that the medical staff as well as the patient can understand the outputs the ASCAPE platform produces. |
| Rationale | Supporting HP.1, PT.1 |
| Supporting materials | N/A. |

Functional Requirement FUNC07

| | |
|---|---|
| ID | FUNC07 |

| | |
|---|---|
| Name | Retrieval of ASCAPE Signals for all patient |
| Priority of accomplishment | **Should have** |
| Description | The ASCAPE-powered Healthcare Provider Information System must be able to retrieve patient-related ASCAPE signals in reverse chronological order and with pagination |
| Rationale | Supporting HP.1 |
| Supporting materials | N/A. |

With regards to the patient-centric data collection use case which aimed to show the potential diversity of data that may be collected about a patient, two things are clear:

- That all sensitive data are collected and connected to the patient's record by the Healthcare Provider Information System and the exact details of their implementation of this data gathering functionality is not within the scope of the present requirements
- That in order for ASCAPE's models to be able to gather data from different ASCAPE-powered Information Systems, there must be a degree of standardisation and a clear understanding of the semantics of data used to train the ASCAPE models. These issues will be addressed further down the line in the project in the dedicated tasks, but here it is obvious that there are requirements related to patient-centric data sources.

Functional Requirement FUNC08

| | |
|---|---|
| ID | FUNC08 |
| Name | Consideration of patient-centric data from wearable devices |
| Priority of accomplishment | **Should have** |
| Description | The ASCAPE AI infrastructure should be able to make use of anonymised wearable-device derived data sent to it by a Healthcare Provider Information System in the context of both training and using AI models |
| Rationale | Supporting HP.1, HP.2, PT.1 |
| Supporting materials | N/A. |

Functional Requirement FUNC09

| | |
|---|---|
| ID | FUNC09 |
| Name | Consideration of patient-centric data from mobile devices |
| Priority of accomplishment | **Should have** |
| Description | The ASCAPE AI infrastructure should be able to make use of anonymised mobile-device derived data (activity data, mini questionnaire) sent to it by a Healthcare Provider Information System in the context of both training and using AI models |
| Rationale | Supporting HP.1, HP.2, PT.1 |
| Supporting materials | N/A. |

Functional Requirement FUNC10

| ID | FUNC10 |
|---|---|
| Name | Consideration of patient-centric data from questionnaires |
| Priority of accomplishment | **Must have** |
| Description | The ASCAPE AI infrastructure should be able to make use of anonymised questionnaire derived data sent to it by a Healthcare Provider Information System in the context of both training and using AI models |
| Rationale | Supporting HP.1, HP.2, PT.1 |
| Supporting materials | N/A. |

Finally, the use case concerning the ASCAPE-specific administration functionality straightforwardly yields a directly relevant functional requirement.

Functional Requirement FUNC11

| ID | FUNC11 |
|---|---|
| Name | Configuring whether or not the ASCAPE-powered Healthcare Provider Information System will be contributing to the ASCAPE global knowledge and how |
| Priority of accomplishment | **Should have** |
| Description | The Administrator must be able to configure whether or not the ASCAPE-powered Healthcare Provider Information System will be contributing to the ASCAPE global knowledge. The default choice must be that it should not. If it is chosen that it does, then the Administrator must be able to specify the method. The two foreseen options which must be provided are via Homomorphic Encryption and via Federated Learning. |
| Rationale | Supporting ADM.1 |
| Supporting materials | N/A. |

## 6.2 Non-functional Requirements

This section gives an overview regarding the non-functional requirements for the system. These are requirements on aspects of quality of the ASCAPE framework that will be essential in it satisfying the user requirements and help make ASCAPE integration an attractive value proposition for healthcare IT system providers.

### 6.2.1 Security Requirements

ASCAPE aims to provide a technological offering comprising local and cloud-based components that will be easy to integrate into existing healthcare IT solutions and be offered by providers of such solutions as an attractive technologically advanced upgrade. This aim can only be achieved if ASCAPE meets or exceeds the security

and privacy standards of existing solutions. General security requirements are addressed first, setting the overall framework in which privacy concerns are to be addressed next (Section 6.2.2).

NonFuncS01 - Authentication, role-based security and access control

| ID | NonFuncS01 |
|---|---|
| Name | Authentication, Authorisation, and Accounting |
| Priority of accomplishment | **Must have** |
| Description | ASCAPE must support:<br>• trustworthy mechanisms for the authentication of third-party entities,<br>• trustworthy mechanisms for the authorisation of entities and the enforcement of access control policies<br>• trustworthy mechanisms for accounting i.e. keeping audit logs of actions and usage of resources |
| Rationale | ASCAPE needs to know the identity of entities that attempt actions, to categorise them by means of a roles system in order to systematise access control rules, to have a system that enforces those access rules and a system for recording access to resources and various parts of the system at a reasonable and suitable granularity level.<br><br>Supporting HP.1, HP.2, PT.1, ADM.1. |
| Supporting materials | N/A |

NonFuncS02 - Integrity

| ID | NonFuncS02 |
|---|---|
| Name | Integrity |
| Priority of accomplishment | **Must have** |
| Description | The integrity of ASCAPE's executables, configuration files and data (including models) must be protected at rest and in transit; only authorised entities (users, software components etc.) must be allowed to make changes and only in the ways allowed. |
| Rationale | Unless the integrity of ASCAPE executables can be guaranteed, ASCAPE cannot be trusted to behave in accordance with its specifications and to meet any other security or other requirement. If its configuration files can be tampered with, it may also behave in undesirable ways, |

| | |
|---|---|
| | possibly even creating significant security loopholes. Finally, if its data can be tampered with, a number of undesirable consequences could ensue including but not limited to, false information about patients being displayed to doctors.<br><br>The requirement does not prescribe the technological means by which it will be met; these will be determined as part of the detailed design of the ASCAPE framework (cloud and edge node). These may include among others: basic (e.g. privilege levels) and advanced (e.g. SGX) CPU-based and OS-based access protection measures for executables and data in the filesystem and in RAM, secure hashing for integrity verification, digital signing and verification, and even measures not associated with security but essential in preserving integrity such as atomicity enforcement measures in ASCAPE business logic implementations and persistence storage access (to avoid integrity being compromised by authorised entities accidentally interfering with one another while attempting to modify the same resource).<br><br>An essential prerequisite for the requirement to be non-vacuous or partial is for an exhaustive catalogue of ASCAPE assets, how they are allowed to be modified and under what conditions (including by whom).<br><br>Supporting HP.1, HP.2, PT.1, ADM.1. |
| Supporting materials | N/A |

NonFuncS03 - Confidentiality

| | |
|---|---|
| ID | NonFuncS03 |
| Name | Confidentiality |
| Priority of accomplishment | **Must have** |
| Description | The confidentiality constraints for any and all ASCAPE assets or parts thereof, such as ASCAPE configuration files and data (including models) must be respected and enforced with appropriate technical means at rest and in transit. |
| Rationale | In a complex system such as ASCAPE overall security relies on secret keys. Additionally, ASCAPE components at the healthcare provider's site process sensitive patient data and must protect them from read access. Even parts |

| | |
|---|---|
| | of code may need to be protected, for IPR or other reasons.<br><br>As noted also in the Rationale of the Integrity requirement, the exact enforcement mechanisms (access control, private-key or public-key cryptography etc.) will be determined as part of ASCAPE's design; moreover, a catalogue of assets with specific read-access constraints will be necessary to make this requirement concrete and enforceable.<br><br>Supporting HP.1, HP.2, PT.1, ADM.1. |
| Supporting materials | N/A |

NonFuncS04 - Availability

| | |
|---|---|
| ID | NonFuncS04 |
| Name | Availability |
| Priority of accomplishment | **Must have** |
| Description | The ASCAPE framework should remain operational under adverse conditions and must protect its ability to recover operation even under extreme conditions. |
| Rationale | Meeting this generic requirement involves preparing to face a number of issues, ranging from hardware failure to malicious actions targeting the integrity of the system and its data or its ability to service legitimate requests. Non-disruption of service is important, but the ability to restore service is crucial.<br><br>The main means of satisfying this requirement is redundancy (e.g. local backups and remote backups) and either over-provisioning or highly-efficient on-demand provisioning to avoid non-malicious threats to availability due to demand peaks, as well as specialised additional measures to thwart malicious distributed denial-of-service (DDOS) attacks.<br><br>Supporting HP.1, HP.2, PT.1, ADM.1. |
| Supporting materials | N/A |

NonFuncS05 - Breach Detection

| | |
|---|---|
| ID | NonFuncS05 |
| Name | Breach Detection |

| | |
|---|---|
| Priority of accomplishment | **Must have** |
| Description | Any detected security breaches and failures of the system to operate in the prescribed manner, must be recorded. Reasonable effort must be made by the ASCAPE framework to detect such events, notify the relevant system administrators (depending on whether the issue is with the ASCAPE Cloud or a local one) and attempt to self-heal where appropriate. |
| Rationale | Not all threats and issues can be avoided, but at least when they are discovered, there is a chance, depending on the kind of incident, to minimise their effect, remedy the problem, detect the conditions under which it occurred, and attempt to avoid it reoccurring.<br><br>Supporting HP.1, HP.2, PT.1, ADM.1. |
| Supporting materials | N/A |

### 6.2.2 Privacy Requirements

Separately addressing privacy concerns serves to emphasise their importance and highlight the relevant ASCAPE technological advantages which allow it to go above and beyond what might be considered possible from a healthcare IT system that aims to collect knowledge from healthcare providers around the world.
Privacy preservation of data is a crucial requirement in any system and especially in systems that analyse and process personal patients' and other healthcare data. Therefore, patient information shall be treated in a highly confidential manner, hence promoting and maintaining fundamental medical ethical principles. Once sensitive information about an individual is exposed, it cannot be withdrawn and made secret again, leading to irreparable damages. Thus, issues on privacy preservation are of major importance for ASCAPE and will be carefully maintained in all levels of the framework.
Two key scenarios will be employed in ASCAPE framework dealing with sensitive patient data (data collected from different sources like healthcare records, Quality of Life data, data collected from patient's wearables, etc.). One scenario involves model inferences being performed locally and federated learning. The second scenario is based on homomorphic encryption which allows sending patient private data to the Cloud infrastructure to be processed remotely without revealing these data to the remote processing site, namely the ASCAPE Cloud.
ASCAPE framework should preserve strong privacy constraints in both scenarios. In the first scenario all crucial activities on data processing happen in the edge-node, without the need to share private data with other parties; only models obtained from the local data, not the data themselves are propagated to the Cloud. The framework should ensure privacy of the edge nodes' personal data, protecting them from unauthorized access. In the second scenario, the framework should also ensure that

the privacy of the sensitive patients' data is not compromised through the analytics produced at the Cloud. Finally, the transfer of the data from an edge node to the cloud should also be secured, preventing malicious data breaches. To completely support the privacy issue of such activity in the framework highly reliable privacy mechanisms will be used, based on Homomorphic Encryption, and Differential privacy.
ASCAPE framework is a complex ecosystem with a large number of software components, hosted and being part of either the ASCAPE Cloud or a Healthcare Provider's ASCAPE-powered IT system. Each component and service that is not directly based on and connected to patient data will implement adequate mechanisms of privacy protection specific for that particular component/service.
The ASCAPE framework should also be adequately aligned with the General Data Protection Regulation (EU) 679/2016 (GDPR), which entered into force on 25 May 2018 and is described in D7.2 "Protection of Personal Data". Chapter 3 of GDPR (Art.12-23) presents the various rights of data subjects, which are summarized to right to information, access, rectification, erasure, restriction, data portability, objection and objection to automated processing.
In this section we addressed insights only into the general privacy requirements as the final ASCAPE architecture is not completely defined.

NonFuncP01 - Patient data privacy inside an edge node

| ID | NonFuncP01 |
|---|---|
| Name | Patient data privacy inside an edge node |
| Priority of accomplishment | **Must have** |
| Description | ASCAPE-powered Healthcare Provider Information System capable to process data locally must ensure privacy of the patients' personal data, protecting it from unauthorized access. |
| Rationale | Elicited from SOTA section 4.2<br>Supporting HP.1, HP.2, PT.1 |
| Supporting materials | N/A. |

NonFuncP02 - Privacy in interaction with ASCAPE Cloud

| ID | NonFuncP02 |
|---|---|
| Name | Privacy in interaction with ASCAPE Cloud |
| Priority of accomplishment | **Must have** |
| Description | ASCAPE-powered Healthcare Provider Information Systems with limited processing resources send anonymised patient data to the Cloud. The security block must ensure privacy preservation of the patient data using highly reliable privacy mechanisms. |
| Rationale | Elicited from SOTA sections 4.3 and 4.4<br>Supporting HP.1, HP.2, PT.1 |
| Supporting materials | N/A. |

NonFuncP03 - Privacy in remote collection of patient data

| ID | NonFuncP03 |
|---|---|

| | |
|---|---|
| Name | Privacy in remote collection of patient data |
| Priority of accomplishment | **Must have** |
| Description | ASCAPE-powered Healthcare Provider Information System must ensure privacy protection of remotely collected patient-centric data (wearables, mobile devices, questionnaires). Remotely collected data must be anonymized before processed further by ASCAPE platform. |
| Rationale | Supporting HP.2, PT.1 |
| Supporting materials | N/A. |

NonFuncP04 - Privacy within instant alerts

| | |
|---|---|
| ID | NonFuncP04 |
| Name | Privacy within instant alerts |
| Priority of accomplishment | **Could have** |
| Description | ASCAPE-powered Healthcare Provider Information System must ensure that instant alerts sent via email do not disclose any sensitive data to third parties. |
| Rationale | Supporting FUNC04 and HP.2 |
| Supporting materials | N/A. |

NonFuncP05 - Privacy in training and updating federated machine learning models

| | |
|---|---|
| ID | NonFuncP05 |
| Name | Privacy in training and updating federated machine learning models |
| Priority of accomplishment | **Must have** |
| Description | All the machine learning models in the ASCAPE framework must be trained and updated in a privacy preserving manner. Special attention regarding this matter must be given to the following tasks: 1) training/updating the federated clients' models with their local data, 2) reducing the received federated clients' models on the federated learning server. It is important that sensitive patient data cannot be discovered or reconstructed from models, i.e. models must not be vulnerable to model inversion attacks. |
| Rationale | If a model is vulnerable to model inversion attacks, private patient data from one edge node can be revealed at another, causing a data breach.<br>Supporting PT.1 |
| Supporting materials | N/A. |

NonFuncP06 - Privacy in inclusion of a new federated partner

| | |
|---|---|
| ID | NonFuncP06 |
| Name | Privacy in inclusion of a new federated partner |

| | |
|---|---|
| Priority of accomplishment | **Must have** |
| Description | Privacy must be guaranteed to a newly registered federated partner. New models, which are initially trained on new federated partners' premises and on their own local data, must be incorporated in the global models in a privacy preserving manner. |
| Rationale | Supporting ADM.1, Requiring NonFuncP05 |
| Supporting materials | N/A. |

NonFuncP07 - Privacy in components/services of Healthcare Provider Information System supporting ASCAPE framework functioning

| | |
|---|---|
| ID | NonFuncP07 |
| Name | Privacy in components/services of Healthcare Provider Information System supporting ASCAPE framework functioning |
| Priority of accomplishment | **Should have** |
| Description | Each component and service that is not directly based on and connected to patients' data will implement adequate mechanisms of privacy protection specific for that particular component/service. Local privacy protocols of Healthcare Provider Information System should be followed. |
| Rationale | |
| Supporting materials | N/A. |

NonFuncP08 - GDPR compliancy

| | |
|---|---|
| ID | NonFuncP08 |
| Name | GDPR compliancy |
| Priority of accomplishment | **Must have** |
| Description | ASCAPE platform must be compliant with the EU GDPR, as well as the national data protection, privacy and ethical legislation in each participant country |
| Rationale | Elicited from D7.2 "Protection of Personal Data" Supporting HP.1, HP.2, PT.1 |
| Supporting materials | [107] |

### 6.2.3 Performance Requirements

The overarching aim set by the following performance requirements is that ASCAPE's performance should not hinder the usability and overall quality of the system. The requirements set forth relate to user-observable aspects of ASCAPE's behaviour, not to operations that might be required to achieve that behaviour.

Despite focusing on user experience rather than the operations upon which ASCAPE relies to provide its functionality, they will be crucial in deciding how to design the ASCAPE infrastructure; this is because, given current technological limitations they can only realistically be met by appropriate design decisions compatible with ASCAPE's aims, not mere brute force.
There are two key user-observable performance-related aspects of ASCAPE's behaviour identified during the requirements elicitation process:

- Response times for ASCAPE-powered results
- Freshness of ASCAPE-powered results

Both focus on the provision of results. Whereas querying a DL model to get a model-predicted data point is a computationally light operation, in order to provide the desired patient-status overview and help achieve ASCAPE-powered results contain and/or are based on multiple such model-predicted data points, so performance issues cannot be ignored when interrogating the models.
These two aspects are addressed by corresponding performance requirements below.
Particular importance is paid to response times for getting an overview of the status of a specific patient (Use Case HP.1 – Patient Visit, Step 4) and common subsequent actions, the requirement being that system response should be practically instantaneous.
Given the combination of the complexity of the computations and the requirement that performance targets are to be achievable without exorbitant hardware infrastructure costs for either the ASCAPE Cloud or the ASCAPE Edge Nodes, it impossible to guarantee the computation of the results will be "practically instantaneous".
A reasonable strategy may be to move the point of the computation to when new data is ingested and caching the result, rather than when a doctor requests to see ASCAPE results about a patient; if this strategy is followed up, the Doctor will be able to get, almost instantly, pre-computed results about the patient (and be notified if there are more recent data that have not been factored into those results). The requirement concerning freshness of ASCAPE-powered results, sets a target for how long the delay between data is ingested and the pre-computed results to be presented to doctors are updated.

NonFuncPf01 - ASCAPE Patient Results Performance

| ID | NonFuncPf01 |
|---|---|
| Name | ASCAPE Patient Results Performance |
| Priority of accomplishment | **Must have** |
| Description | Response times for ASCAPE to provide the patients overview results page of Use Case HP.1 – Patient Visit, Step 4 must be less than 3 seconds.<br><br>Responses to most common subsequent user actions must also be less than 3 seconds; for the remaining actions appropriate indication that more time will be required must be displayed and results should be returned in a time commensurate with the complexity of the task. |

| | |
|---|---|
| Rationale | The ASCAPE-powered visualisation and simulation component integrated into a healthcare provider system of the edge node needs to be highly responsive and not insert noticeable delays in the doctors' interaction with the healthcare provider's system when the doctor views the patient's records, including the ASCAPE-provided patient status overview described in Step 4 of the Use Case HP.1 – Patient Visit.<br>Moreover, commonly requested ASCAPE-provided results must be available instantly. Such a strict requirement for instantaneous responses supports the usability requirements with respect to Use Case HP.1 – Patient Visit.<br><br>On the other hand, the requirement avoids eliminating possibly desirable functionality that cannot optimised by means of pre-computation. If the detailed ASCAPE design for the application of ASCAPE in a specific field (e.g. breast or prostate cancer patients' QoL) leads to a conviction that a certain functionality that requires on-the-spot computation is beneficial, the requirement allows it to be included. User Acceptance Testing, performance optimisation experimentation and a broader assessment of the usefulness of the feature, not a single-dimensional performance requirement, will determine if it will be included after all.<br>Supporting HP.1 |
| Supporting materials | N/A. |

NonFuncPf02 - ASCAPE Patient Data Processing Delays

| | |
|---|---|
| ID | NonFuncPf02 |
| Name | ASCAPE Patient Data Processing Delays |
| Priority of accomplishment | **Should have** |
| Description | ASCAPE outputs based on questionnaires and any user-provided data should be available to doctors via ASCAPE in less than 2 minutes from the time they are entered in the ASCAPE Edge Node-connected system where the data are collected.<br><br>ASCAPE outputs based on data from wearables or other tracking devices (incl. mobiles) up to midnight of a day should be available to doctors via ASCAPE by 8 am of the following day. |
| Rationale | In most cases significantly more lax requirements would suffice. For example, in a situation where patients fill in questionnaires the day prior to their doctor visit, ASCAPE |

| | |
|---|---|
| | would have hours to process the freshly provided data and pre-compute and cache ASCAPE results for that patient. Likewise, as the results of blood tests and other diagnostic examinations are delivered with a delay and doctor's appointments are scheduled with that reality in mind, there typically is more than enough time for ASCAPE to produce its results on the basis of their outcomes ahead of a doctor's appointment.<br><br>However, in order for ASCAPE to not come with constraints that may not be compatible with various healthcare providers' protocols, the requirement sets a target for a delay of no more than 2 minutes, which is a sufficiently short period of time to alleviate this concern in most cases. For example, if according to a healthcare provider's protocol, the patient fills in a questionnaire in the waiting area before they see their doctor, it is more than likely that ASCAPE will have enough time to precompute and cache results before the patient enters the doctor's office and the doctor accesses their record.<br><br>Finally, the requirement does not preclude designs that make an effort to further improve user experience in cases where data are added during the doctor visit (or right before it); this may involve on-the-spot gradual updates of the ASCAPE-results.<br><br>Supporting HP.1 |
| Supporting materials | N/A. |

There are no performance requirements regarding freshness of models used for providing ASCAPE-powered results as for privacy reasons (thwarting privacy attacks based on comparing predictions of the current and the previous version of a model in order to obtain information about newly added data), it is preferable to update models infrequently on batches of patient data.

### 6.2.4 Hardware Support Requirements

ASCAPE aims both to build a common cloud infrastructure and to provide an easy migration path for software vendors to provide ASCAPE-powered functionality in their information systems. In order to facilitate both the creation (and later expansion) of the ASCAPE cloud and the insertion of ASCAPE edge nodes in existing healthcare IT infrastructures, a number of requirements about hardware support are provided below, including requirements that ensure that certain hardware features, if present, are taken advantage of.

#### 6.2.4.1 ASCAPE Cloud Infrastructure

Given the fact that the ASCAPE Cloud may have its entire hardware infrastructure upgraded without systems connected to it needing to be aware of the details of its (hardware and software) implementation, it is not necessary to be overly specific about hardware requirements for the nodes that will comprise the ASCAPE Cloud.
The only two requirements concerning the ASCAPE Cloud infrastructure are meant to direct development of the ASCAPE Cloud infrastructure towards the hardware architectures that are most commonly used today. The idea is to ensure the ASCAPE Cloud can expand on existing or newly purchased servers and take advantage of a market leading series of General Purpose Graphics Processing Units (GPGPUs) with good toolkit support for Deep Learning.

Non-FuncH01 - ASCAPE cloud x86-64 CPUs support

| ID | Non-FuncH01 |
|---|---|
| Name | ASCAPE cloud x86-64 CPUs support |
| Priority of accomplishment | **Must have** |
| Description | ASCAPE cloud components must be capable of running on x86-64 instruction set CPUs |
| Rationale | The most widely-used server CPUs currently and in the foreseeable future are AMD and Intel 64-bit CPUs supporting the x86-64 instruction set. This requirement does not exclude support for other CPUs either now or in the future but ensures the most widely used ones are supported. |
| Supporting materials | N/A. |

Non-FuncH02 - ASCAPE cloud GPGPUs support

| ID | NonFuncH02 |
|---|---|
| Name | ASCAPE cloud GPGPUs support |
| Priority of accomplishment | **Should have** |
| Description | When GPGPUs are available on ASCAPE Cloud servers, ASCAPE should be able to utilise them to provide enhanced performance |
| Rationale | GPGPUs can assist in executing computationally intensive tasks, like complex simulations based on ASCAPE DL models and DL model updates etc. offering a significant improvement in performance. A decision has not been made about which specific GPGPUs are to be supported.<br>Supporting HP.1 and HP.2 |
| Supporting materials | N/A. |

#### 6.2.4.2 Edge node

One way a software vendor may provide ASCAPE-powered functionality in their information system is by re-implementing or integrating at the source-code level ASCAPE non-cloud components which provide the functionality of ASCAPE's Edge Node software, as provided by the ASCAPE Consortium. In that case, a physical or a virtual machine or possibly even a container will be required to host the ASCAPE Edge Node software.
In order to make the choice to use an ASCAPE Edge Node easy to support, minimum requirements need to be set; the onus will be on the ASCAPE architecture, on the ASCAPE technologies and on ASCAPE code performance tuning to ensure the performance requirements are met even with a bare minimum edge-node infrastructure. On the other hand, where additional resources are available, they are to be utilised for performance improvements.

Non-FuncH03 - ASCAPE edge node x86-64 CPUs support

| | |
|---|---|
| ID | Non-FuncH03 |
| Name | ASCAPE edge node x86-64 CPUs support |
| Priority of accomplishment | **Must have** |
| Description | ASCAPE edge components must be capable of running on x86-64 instruction set CPUs |
| Rationale | The most widely-used server CPUs currently and in the foreseeable future are AMD and Intel 64-bit CPUs supporting the x86-64 instruction set. This requirement does not exclude support for other CPUs either now or in the future but ensures the most widely used ones are supported.<br>Supporting HP.1 and HP.2 |
| Supporting materials | N/A. |

Non-FuncH04 - Minimum processing capabilities of ASCAPE edge node

| | |
|---|---|
| ID | Non-FuncH04 |
| Name | Minimum processing capabilities of ASCAPE edge node |
| Priority of accomplishment | **Must have** |
| Description | An ASCAPE Edge Node should not require more than the equivalent of a sixth generation i5 processor to meet performance requirements for the core user interactions |
| Rationale | An ASCAPE edge node will need computing power to run simulations, interrogating and training the AI models, performing homomorphic encryption/ decryption, etc. However, by means of clever design the core user interactions can be served by pre-computed results; this will not be possible for secondary interactions such as the exploration of different treatment scenarios. Additionally, the ASCAPE Edge Node will offer technological choices |

| | |
|---|---|
| | such as homomorphic encryption which can help offload tasks to the ASCAPE Cloud.<br>Supporting HP.1 and HP.2 |
| Supporting materials | N/A. |

Non-FuncH05 - Minimum memory requirements for ASCAPE edge node

| | |
|---|---|
| ID | Non-FuncH05 |
| Name | Minimum memory requirements for ASCAPE edge node |
| Priority of accomplishment | **Must have** |
| Description | An ASCAPE Edge Node should be capable of operating with modest memory resources (minimum: 4GB DDR4 RAM) |
| Rationale | Memory is a relative expensive resource and minimising memory requirements will help ASCAPE adoption. This will be possible given the fact that the ASCAPE Edge Node will offer technological choices such as homomorphic encryption which can help offload tasks to the ASCAPE Cloud.<br>Supporting HP.1 and HP.2 |
| Supporting materials | N/A. |

Non-FuncH06 - Minimum storage requirements for ASCAPE edge node

| | |
|---|---|
| ID | Non-FuncH06 |
| Name | Minimum storage requirements for ASCAPE edge node |
| Priority of accomplishment | **Must have** |
| Description | An ASCAPE Edge Node should be capable of operating with modest persistent storage capacity (minimum: 30GB) |
| Rationale | The ASCAPE Edge Node does not need to store large quantities of data, but rather to act as an intermediary enabling the collaboration between the healthcare system where patient data are stored and the ASCAPE Cloud where global models reside, but it does needs storage for ASCAPE Edge components and supporting software infrastructure, local models and local copies of global models, configuration files, logs etc.<br>Supporting HP.1 and HP.2 |
| Supporting materials | N/A. |

Non-FuncH07 - Minimum network utilisation by ASCAPE edge node and off-line operation

| | |
|---|---|
| ID | Non-FuncH07 |
| Name | Minimum network utilisation by ASCAPE edge node and off-line operation |
| Priority of accomplishment | **Must have** |

| | |
|---|---|
| Description | The ASCAPE Edge Node must not require a connection to the ASCAPE Cloud with a bandwidth of over 1Mbs and roundtrip latency (ping) of less than 100ms and a connection with the healthcare provider's system with a bandwidth of over 10Mbs and roundtrip latency (ping) of less than 100ms. |
| Rationale | This requirement essentially provides reassurance that no special network provisions will be required for the ASCAPE Edge Node, covering also the case of healthcare providers that may not desire to have the ASCAPE Edge Node connecting to the Internet and the ASCAPE Cloud in particular. In the standard scenario, the ASCAPE Edge Node will create only reasonably small amounts of traffic when communicating with the ASCAPE Cloud (internet connection) and the healthcare provider's system (which would normally be over a local connection), thus not overwhelming the respective connections even with if they have the minimum bandwidth required.<br>Supporting HP.1 and HP.2 |
| Supporting materials | N/A. |

Non-FuncH08 - ASCAPE edge node GPGPUs support

| | |
|---|---|
| ID | NonFuncH08 |
| Name | ASCAPE edge node GPGPUs support |
| Priority of accomplishment | **Should have** |
| Description | When GPGPUs are available for ASCAPE Edge Nodes to use, they should be able to utilise them to provide enhanced performance |
| Rationale | GPGPUs can assist in executing computationally intensive tasks, like complex simulations based on ASCAPE DL models and DL model updates etc. offering a significant improvement in performance. A decision has not been made about which specific GPGPUs are to be supported.<br>Supporting HP.1 and HP.2 |
| Supporting materials | N/A. |

### 6.2.5 Usability

Usability will be a key factor in the success of ASCAPE; doctors need not only be convinced that ASCAPE can provide useful results, but also feel comfortable using the relevant ASCAPE-powered UI. This UI will need to be cleverly integrated into the UI of the information systems they currently use. These requirements apply to the ASCAPE Dashboard (D4.1) where the ASCAPE UI elements will first be incorporated and demonstrated but should also be taken into consideration when creating

ASCAPE-powered versions of healthcare provider IT systems (as will be demonstrated also in WP4).

Non-FuncU01 - Learnability

| | |
|---|---|
| ID | Non-FuncU01 |
| Name | Learnability |
| Priority of accomplishment | Should have |
| Description | The system should be easy to learn for both inexperienced and experienced users. |
| Rationale | Supporting HP.1 and HP.2 |
| Supporting materials | N/A |

Non-FuncU02 - Memorability

| | |
|---|---|
| ID | Non-FuncU02 |
| Name | Memorability |
| Priority of accomplishment | Should have |
| Description | The system should be easy to remember for the casual user. |
| Rationale | Supporting HP.1, HP.2 and ADM.1 |
| Supporting materials | N/A |

Non-FuncU03 - Error feedback and recovery

| | |
|---|---|
| ID | Non-FuncU03 |
| Name | Error feedback and recovery |
| Priority of accomplishment | Should have |
| Description | Any errors (of input or otherwise) will be communicated to the user in a straightforward manner, and their impact will be clear, and if applicable, recoverable. |
| Rationale | Supporting HP.1 and ADM.1 |
| Supporting materials | N/A |

Non-FuncU04 - Satisfaction

| | |
|---|---|
| ID | Non-FuncU04 |
| Name | Satisfaction |
| Priority of accomplishment | Must have |
| Description | The user will feel satisfied with the use of the system and will be more likely to recommend it to other patients, than not. |
| Rationale | Supporting HP.1, HP.2, PT.1 and ADM.1 |
| Supporting materials | N/A |

Non-FuncU05 - Consistent navigation

| | |
|---|---|
| ID | Non-FuncU05 |
| Name | Consistent navigation |

| | |
|---|---|
| Priority of accomplishment | Should have |
| Description | The navigation and content structure must be coherent throughout the system. To this end, i) The same action should produce always the same response, ii) Links, action buttons and objects must be organized coherently, iii) High importance messages should be visible upon login, iv) Response times should be appropriated for each task, and v) Data entries should not be case sensitive and should clearly state which kind of data do they accept |
| Rationale | Supporting HP.1 and HP.2 |
| Supporting materials | N/A |

Non-FuncU06 - Task efficiency

| | |
|---|---|
| ID | Non-FuncU06 |
| Name | Task efficiency |
| Priority of accomplishment | Should have |
| Description | During routine appointments, the doctor will be able to view the patient's journey regarding their physical, social, psychological and other consequences from the condition and its treatments. The system must be efficient for the frequent user. |
| Rationale | Supporting HP.1 and HP.2 |
| Supporting materials | N/A |

Non-FuncU07 - Clear organisation of information

| | |
|---|---|
| ID | Non-FuncU07 |
| Name | Clear organisation of information |
| Priority of accomplishment | Should have |
| Description | All the information on the system must be well organised. |
| Rationale | Supporting HP.1 and HP.2 |
| Supporting materials | N/A |

### 6.2.6 Overall Quality Requirements

Non-FuncQ01 - State of the art analytics

| | |
|---|---|
| ID | NonFuncQ01 |
| Name | Adaptability |
| Priority of accomplishment | **Should have** |
| Description | The ASCAPE platform should be able to be easily adjusted to cater to different data schemata and formats and other types of cancer |
| Rationale | This requirement reinforces the ASCAPE ambition of being able to be adjusted to different environments and |

| | |
|---|---|
| | needs, including future support for different cancer types. ASCAPE's Adaptability will be tested in practice during the Open Call.<br>Supporting HP.1, HP.2 and PT.1 |
| Supporting materials | [108] |

Non-FuncQ02 - Functional and flexible operation

| | |
|---|---|
| ID | NonFuncQ02 |
| Name | Functional and flexible operation |
| Priority of accomplishment | **Should have** |
| Description | The ASCAPE platform must be able to support the functional, flexible and efficient operation in a distributed cloud infrastructure |
| Rationale | The operational compatibility of the platform of the ASCAPE platform is a crucial quality characteristic for the platform's operation and reusability.<br>Supporting HP.1, HP.2, PT.1 and ADM.1 |
| Supporting materials | [108] |

Non-FuncQ03 - Interoperability

| | |
|---|---|
| ID | NonFuncQ03 |
| Name | Interoperability |
| Priority of accomplishment | **Should have** |
| Description | The ASCAPE platform should be able to support the interconnection and exchange of information with other platforms/devices in a secure manner. |
| Rationale | The interoperability of the platform of the ASCAPE platform is a crucial quality characteristic for the platform's compatibility, extensibility and exploitation potentials.<br>Supporting HP.1, HP.2, PT.1 and ADM.1 |
| Supporting materials | [108] |

Non-FuncQ04 - High availability

| | |
|---|---|
| ID | NonFuncQ04 |
| Name | High availability |
| Priority of accomplishment | **Must have** |
| Description | The ASCAPE platform should be able to ensure high availability of the system and the stored information. |
| Rationale | The high availability of the ASCAPE platform is a crucial quality characteristic for the platform's reliability.<br>Supporting HP.1, HP.2, PT.1 and ADM.1 |
| Supporting materials | [108] |

Non-FuncQ05 - Recovery and Fault-tolerance

| | |
|---|---|
| ID | NonFuncQ05 |
| Name | Recovery and Fault-tolerance |
| Priority of accomplishment | **Must have** |
| Description | The ASCAPE platform must be able to able to recover from system failure conditions and effectively handle software failure condition without affecting the platform's overall functional operation. |
| Rationale | The recoverability and fault-tolerance of the ASCAPE platform is a crucial quality characteristic for the platform's reliability.<br>Supporting HP.1, HP.2, PT.1 and ADM.1 |
| Supporting materials | [108] |

Non-FuncQ06 - Portability

| | |
|---|---|
| ID | NonFuncQ06 |
| Name | Modularity |
| Priority of accomplishment | **Should have** |
| Description | The ASCAPE platform should be composed of independent components that have well defined interfaces and are replaceable with minimum impact and effort |
| Rationale | The evolution of ASCAPE efficient deployment of the ASCAPE platform, as well as the efficient replacement of the components of the platform if needed is a crucial quality characteristic for the platform's portability.<br>Supporting HP.1, HP.2, PT.1 and ADM.1 |
| Supporting materials | [108] |

## 6.3 Relation to State of the Art Advancements

The above requirements are strongly related also to the State of the Art advancements being pursued in ASCAPE (see Section 3). The following table captures these relations.

| SotA Advancement | Requirement | Relation |
|---|---|---|
| Explainable AI for Healthcare | Functional Requirement FUNC03 - Generation of ASCAPE Signals for a given patient<br>Functional Requirement FUNC04 - Instant alert upon the generation of ASCAPE Signals for a given patient | Once doctors select a signal on the patient status overview dashboard or investigates a signal-caused alert, they are presented with a series of graphs, ordered by importance, which provide a |

| SotA Advancement | Requirement | Relation |
|---|---|---|
| | | quick overview of the status of the patient's QoL and provide visual support for the signal. |
| | Functional Requirement FUNC01 - Generation of ASCAPE What-If Graph for a given patient, a given metric and a (possibly empty) set of potential QoL interventions<br>Functional Requirement FUNC02 - Calculation of a prominence metric for each ASCAPE What-If Graph | The What-If graphs not only give a quick overview of a patient's QoL status or help explain ASCAPE signals, but also provide a means for the doctor to experiment, for example to try the effect of different interventions; thus, they not only attain an understanding of why a particular intervention was chosen by the system but also what the system believes about alternative interventions. |
| | Functional Requirement FUNC05 - All predictions of the models should be explainable<br>Functional Requirement FUNC06 - The medical intervention suggestions of the models should be explainable | In addition to the level of explanation that is possible with What-If graphs, ASCAPE will be able to provide an additional degree of explainability on the basis of relevant State-of-the-Art (Advancement 1 and 3) |
| Federated Learning for Healthcare<br><br>Homomorphic Encryption for Healthcare | Functional Requirement FUNC01 - Generation of ASCAPE What-If Graph for a given patient, a given metric and a (possibly empty) set of potential QoL interventions<br>Functional Requirement FUNC03 - Generation of ASCAPE Signals for a given patient | The core AI functionality of ASCAPE is provided by model inference either on standard DL models or HE-DL models. These are created by the means described in the State of the Art Advancement sections for Federated Learning and Homomorphic Encryption for Healthcare respectively. |

| SotA Advancement | Requirement | Relation |
|---|---|---|
| | Functional Requirement FUNC08 - Consideration of patient-centric data from wearable devices)<br>Functional Requirement FUNC09 - Consideration of patient-centric data from mobile devices)<br>Functional Requirement FUNC10 - Consideration of patient-centric data from questionnaires | The model training functionality will be provided using the two technologies outlined in the relevant sections presenting the ASCAPE advances in Federated Learning for Healthcare and the ASCAPE Advances in Homomorphic Encryption for Healthcare. |
| | Functional Requirement FUNC11 - Configuring whether or not the ASCAPE-powered Healthcare Provider Information System will be contributing to the ASCAPE global knowledge and how | The requirement concerns controlling which of the two available technologies (outlined in the relevant sections presenting the ASCAPE advances in Federated Learning for Healthcare and the ASCAPE Advances in Homomorphic Encryption for Healthcare) or combination thereof is to be used for updating global ASCAPE models using local data. |
| Epsilon-Differential Privacy | Functional Requirement FUNC08 - Consideration of patient-centric data from wearable devices<br>Functional Requirement FUNC09 - Consideration of patient-centric data from mobile devices<br>Functional Requirement FUNC10 - Consideration of patient-centric data from questionnaires | Epsilon-Differential Privacy may be applied to the different types of data ingested by an ASCAPE Edge Node. |

| SotA Advancement | Requirement | Relation |
|---|---|---|
| | NonFuncP01 - Patient data privacy inside an edge node<br>NonFuncP02 - Privacy in interaction with ASCAPE Cloud<br>NonFuncP05 - Privacy in training and updating federated machine learning models<br>NonFuncP06 - Privacy in inclusion of a new federated partnerNonFuncP07 - Privacy in components/services of Healthcare Provider Information System supporting ASCAPE framework functioning | |

Table 1: Relation of State of the Art Advancements to Requirements

# 7 Conclusions

ASCAPE is an ambitious research project, that focuses on the development of user-centric healthcare solutions to improve the quality of life after cancer treatment.
One of the key points of ASCAPE solutions is that patients' medical data are treated with confidentiality and integrity. This is achieved through the use of innovative technologies such as federated machine learning, epsilon differential privacy and homomorphic encryption. The sensitive data are: 1) either processed on the edge node to create local models which are sent to the Cloud or 2) are homomorphically encrypted and then the encrypted data are sent to the Cloud to be processed remotely. In both cases, the general knowledge is updated while privacy is preserved.
It is important to highlight that the ASCAPE solution is designed to be a practical tool for healthcare providers, who will be responsible for determining how to introduce it in their patients' treatment. AI is not mature enough to suggest interventions directly to the patient without doctor's medical opinion. For this reason, ASCAPE's main users are healthcare professionals and its focus is to provide concise and user friendly presentation of results in accordance to their requirements.
Moreover, ASCAPE's exploitation plan aims at creating synergies rather than appearing as competitor to existing healthcare provider systems. This way, the benefits of ASCAPE will spread to all the potential software provider systems that implement ASCAPE-powered versions to their existing software. All those key points of ASCAPE described above, play an important role in delivering the promise of democratization of big data medical knowledge by providing equal access to knowledge even to doctors in areas with low healthcare quality scores (poor countries, remote villages, etc) and by having the capability to address a number of medical issues.
Creating a solution that can support a number of medical issues and many different types of cancer is complicated, as cancer is a group of more than 100 different diseases, so ASCAPE will initiate its work with breast and prostate cancer. The breast and prostate cancer pilots, that will take place on WP4, will provide foundation towards the innovation of democratization of knowledge. The follow-up deliverable D1.2 "ASCAPE Data Determinants and piloting validations", will describe pilot requirements and data determinants and will analyse concrete ways of evaluating ASCAPE in the context of the pilots.
ASCAPE framework requirements that have been specified in this deliverable will set the basis for the architecture design of ASCAPE project that will be presented in the deliverable D1.3 "Architecture definition". Finally, ASCAPE project is well placed to challenge the orthodoxy of the Iron Triangle of Health, as discussed in Section 2.2 and deliverable D1.4 "Manuscripts on costs and benefits of the new diagnostic tool" will present the health economics model that will be used to prove this belief.

# Appendix

The total of ASCAPE requirements are assembled in the following table.

| Requirement ID | Name | Priority of accomplishment | Related to Use cases |
|---|---|---|---|
| **FUNC01** | Generation of ASCAPE What-If Graph for a given patient, a given metric and a (possibly empty) set of potential QoL interventions | **Must have** | HP.1, HP.2 |
| **FUNC02** | Calculation of a prominence metric for each ASCAPE What-If Graph | **Could have** | HP.1, HP.2 |
| **FUNC03** | Generation of ASCAPE Signals for a given patient | **Should have** | HP.1, HP.2 |
| **FUNC04** | Instant alert upon the generation of ASCAPE Signals for a given patient | **Could have** | HP.2 |
| **FUNC05** | All predictions should be explainable | **Must have** | HP.1, HP.2, PT.1 |
| **FUNC06** | The medical intervention suggestions should be explainable | **Must have** | HP.1, PT.1 |
| **FUNC07** | Retrieval of ASCAPE Signals for all patient | **Should have** | HP.1 |
| **FUNC08** | Consideration of patient-centric data from wearable devices | **Should have** | HP.1, HP.2, PT.1 |
| **FUNC09** | Consideration of patient-centric data from mobile devices | **Should have** | HP.1, HP.2, PT.1 |
| **FUNC10** | Consideration of patient-centric data from questionnaires | **Must have** | HP.1, HP.2, PT.1 |
| **FUNC11** | Configuring whether or not the ASCAPE-powered Healthcare Provider Information System will be contributing to the ASCAPE global knowledge and how | **Should have** | ADM.1 |
| **NonFuncS01** | Authentication, role-based security and access control | **Must have** | HP.1, HP.2, PT.1, ADM.1 |
| **NonFuncS02** | Integrity | **Must have** | HP.1, HP.2, PT.1, ADM.1 |
| **NonFuncS03** | Confidentiality | **Must have** | HP.1, HP.2, PT.1, ADM.1 |
| **NonFuncS04** | Availability | **Must have** | HP.1, HP.2, PT.1, ADM.1 |
| **NonFuncS05** | Breach Detection | **Must have** | HP.1, HP.2, PT.1, ADM.1 |
| **NonFuncP01** | Patient data privacy inside an edge node | **Must have** | HP.1, HP.2, PT.1, |
| **NonFuncP02** | Privacy in interaction with ASCAPE Cloud | **Must have** | HP.1, HP.2, PT.1 |

| | | | |
|---|---|---|---|
| **NonFuncP03** | Privacy in remote collection of patient data | **Must have** | HP.2, PT.1 |
| **NonFuncP04** | Privacy within instant alerts | **Could have** | HP.2 |
| **NonFuncP05** | Privacy in training and updating federated machine learning models | **Must have** | PT.1 |
| **NonFuncP06** | Privacy in inclusion of a new federated partner | **Must have** | ADM.1 |
| **NonFuncP07** | Privacy in components/services of Healthcare Provider Information System supporting ASCAPE framework functioning | **Should have** | - |
| **NonFuncP08** | GDPR compliancy | **Must have** | HP.1, HP.2, PT.1 |
| **NonFuncPf01** | ASCAPE Patient Results Performance | **Must have** | HP.1 |
| **NonFuncPf02** | ASCAPE Patient Data Processing Delays | **Should have** | HP.1 |
| **Non-FuncH01** | ASCAPE cloud x86-64 CPUs support | **Must have** | - |
| **Non-FuncH02** | ASCAPE cloud CUDA-enabled GPGPUs support | **Must have** | HP.1, HP.2 |
| **Non-FuncH03** | ASCAPE edge node x86-64 CPUs support | **Must have** | HP.1, HP.2 |
| **Non-FuncH04** | Minimum processing capabilities of ASCAPE edge node | **Must have** | HP.1, HP.2 |
| **Non-FuncH05** | Minimum memory requirements for ASCAPE edge node | **Must have** | HP.1, HP.2 |
| **Non-FuncH06** | Minimum storage requirements for ASCAPE edge node | **Must have** | HP.1, HP.2 |
| **Non-FuncH07** | Minimum network utilisation by ASCAPE edge node and off-line operation | **Must have** | HP.1, HP.2 |
| **Non-FuncH08** | ASCAPE edge node CUDA-enabled GPGPUs support | **Must have** | HP.1, HP.2 |
| **Non-FuncU01** | Learnability | **Must have** | HP.1, HP.2, |
| **Non-FuncU02** | Memorability | **Must have** | HP.1, HP.2, ADM.1 |
| **Non-FuncU03** | Error feedback and recovery | **Should have** | HP.1, ADM.1 |
| **Non-FuncU04** | Satisfaction | **Must have** | HP.1, HP.2, PT.1, ADM.1 |
| **Non-FuncU05** | Consistent navigation | **Should have** | HP.1, HP.2 |
| **Non-FuncU06** | Task efficiency | **Must have** | HP.1, HP.2 |
| **Non-FuncU07** | Clear organization of information | **Must have** | HP.1, HP.2, |
| **NonFuncQ01** | Adaptability | **Should have** | HP.1, HP.2, PT.1 |
| **NonFuncQ02** | Functional and flexible operation | **Should have** | HP.1, HP.2, PT.1, ADM.1 |
| **NonFuncQ03** | Interoperability | **Should have** | HP.1, HP.2, PT.1, ADM.1 |
| **NonFuncQ04** | High availability | **Must have** | HP.1, HP.2, PT.1, ADM.1 |

| | | | |
|---|---|---|---|
| **NonFuncQ05** | Recovery and Fault-tolerance | **Must have** | HP.1, HP.2, PT.1, ADM.1 |
| **NonFuncQ06** | Portability | **Should have** | HP.1, HP.2, PT.1, ADM.1 |

Table 2-Total of ASCAPE requirements

# References

[1] G. Carioli, P. Bertuccio and P. Boffeta, "European cancer mortality predictions for the year 2020 with a focus on prostate cancer," *Ann Oncol.,* no. 31(5), online: https://www.annalsofoncology.org/article/S0923-7534(20)36056-7/fulltext, p. 650-658 doi:10.1016/j.annonc.2020, 2020.

[2] R. Siegel, K. Miller and A. Jemal, "Cancer statistics," *CA Cancer J Clin. 2020,* no. online: https://acsjournals.onlinelibrary.wiley.com/doi/full/10.3322/caac.21590, p. 70(1):7-30. doi:10.3322/caac.21590, 2020.

[3] T. L. Oncology, "Perceptions of cancer in society must change," *The Lancet Oncology,* vol. 17, no. 3 online: https://doi.org/10.1016/S1470-2045(16)00091-7, p. 257, 2016.

[4] M. A. Mayer, A. Heinrich, A. Rodríguez and e. al, "Big Data Technologies in Healthcare. Needs, opportunities and challenges. Big Data Value Association (BDVA)," *10.13140/RG.2.2.35249.89448,* 2016.

[5] J. Ferlay, M. Colombet, I. Soerjomataram, T. Dyba, G. Randi, M. Bettio, A. Gavin, O. Visser and F. Bray, "Cancer incidence and mortality patterns in Europe: Estimates for 40 countries and 25 major cancers in 2018," *Eur J Cancer,* vol. vol. 103, p. 356–387, 2018.

[6] U. Dafni, Z. Tsourti and I. Alatsathianos, "Breast Cancer Statistics in the European Union: Incidence and Survival across European Countries," *Breast Care,* vol. vol. 14, p. 344–352, 2019.

[7] International Agency for Research on Cancer IARC-WHO, "Europe Fact Sheets," 2018. [Online]. Available: https://gco.iarc.fr/today/data/factsheets/populations/908-europe-fact-sheets.pdf.

[8] Crocetti, "Epidemiology of prostate cancer in Europe," 2015. [Online]. Available: https://ec.europa.eu/jrc/en/publication/epidemiology-prostate-cancer-europe.

[9] F. Mols, M. Thong, P. Vissers, T. Nijsten and L. v. d. Poll-Franse, "Socio-economic implications of cancer survivorship: results from the PROFILES registry," *Eur J Cancer,* vol. vol. 48(13), p. 2037–2042, 2012.

[10] T. Albreht, R. Kiasuwa and M. Van den Bulcke, "European Guide on Quality Improvement in Comprehensive Cancer Control," National Institute of Public Health and Scientific Institute of Public health, Ljubljana and Brussels, 2017.

[11] GSMA , " The Mobile Economy Europe 2018 - The Mobile Economy," 2018. [Online]. Available: https://www.gsma.com/mobileeconomy/europe/.

[12] GSMA, "The Mobile Economy 2020," 2020. [Online]. Available: https://www.gsma.com/mobileeconomy/.

[13] Statista, "Wearables sales worldwide by region 2015-2022," 2020. [Online]. Available: https://www.statista.com/statistics/490231/wearable-devices-worldwide-by-region/.

[14] eurostat, ""Population structure and ageing - Statistics Explained," Eurostat, 2020. [Online]. Available: [Accessed: 02-May-2020].," 2020. [Online]. Available: https://ec.europa.eu/eurostat/statistics-explained/index.php/Population_structure_and_ageing.

[15] Y. H. Wu, S. Damnée, H. Kerhervé, C. Ware and A. S. Rigaud, "Bridging the digital divide in older adults: A study from an initiative to inform older adults about new technologies," *Clin. Interv. Aging,* vol. 10, p. 193–201, 2015.

[16] FEBA, "Poverty in Europe," 2018. [Online]. Available: https://www.eurofoodbank.org/en/poverty-in-europe.

[17] European Parliament, "REGULATION (EU) 2017/745 on medical devices," 5 April 2017. [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32017R0745.

[18] W. L. Kissick, Medicine's Dilemmas: Infinite Needs Versus Finite Resources, New Haven: Yale University Press, 1994.

[19] I. Team, "AI And Healthcare: A Giant Opportunity," *Intel AIFORBES INSIGHTS, Issue 06,* pp. Retrieved from: https://www.bibme.org/citation-guide/apa/magazine/, 11 February 2019.

[20] R. Girshick, J. Donahue, T. Darrell and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition,* p. 580–587, 2014.

[21] W. R. Swartout, Rule-based expert systems: The mycin experiments of the stanford heuristic programming project, vol. 26, B. G. Buchanan and E. H. Shortliffe, Eds., Reading, MA: Addison Wesley, 1985, p. 364–366.

[22] A. Holzinger, C. Biemann, C. S. Pattichis and D. B. Kell, "What do we need to build explainable AI systems for the medical domain?," p. 1–28, 2017.

[23] L. Breiman, J. H. Friedman, R. A. Olshen and C. J. Stone, Classification and regression trees, Wadsworth InternationGroup, Belmont, CA, 2017, p. 1–358.

[24] T. A. Etchells and P. J. G. Lisboa, "Orthogonal Search-Based Rule Extraction (OSRE) for Trained Neural Networks: A Practical and Efficient Approach," *IEEE Transactions on Neural Networks,* vol. 17, p. 374–384, 3 2006.

[25] T. Rögnvaldsson, T. A. Etchells, L. You, D. Garwicz, I. Jarman and P. J. G. Lisboa, "How to find simple and accurate rules for viral protease cleavage specificities," *BMC Bioinformatics,* vol. 10, p. 149, 5 2009.

[26] V. Van Belle and P. Lisboa, "White box radial basis function classifiers with component selection for clinical prediction models," *Artificial Intelligence in Medicine,* vol. 60, p. 53–64, 1 2014.

[27] V. Van Belle, B. Van Calster, S. Van Huffel, J. A. K. Suykens and P. Lisboa, "Explaining support vector machines: A color based nomogram," *PLoS ONE,* vol. 11, 10 2016.

[28] C.-A. Peña-Reyes and M. Sipper, “Fuzzy CoCo: Balancing Accuracy and Interpretability of Fuzzy Models by Means of Coevolution,” Springer, Berlin, Heidelberg, 2003, p. 119–146.

[29] M. A. Barreto-Sanz, A. Bujard and C. A. Peña-Reyes, “Evolving very-compact fuzzy models for gene expression data analysis,” in *IEEE 12th International Conference on BioInformatics and BioEngineering, BIBE 2012*, 2012.

[30] J. Despraz, S. Gomez, H. F. Satizábal and C. A. Peña-Reyes, “Exploring internal representations of deep neural networks,” in *Studies in Computational Intelligence*, 2019.

[31] S. Gomez Schnyder, J. Despraz and C. A. Peña-Reyes, “Improving neural network interpretability via rule extraction,” in *Artificial Neural Networks and Machine Learning – ICANN 2018*, Cham, 2018.

[32] M. T. Ribeiro, S. Singh and C. Guestrin, “"Why should i trust you?" Explaining the predictions of any classifier,” *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining,* Vols. 13-17-Augu, p. 1135–1144, 2016.

[33] Z. Zhou, H. Cai, S. Rong, Y. Song, K. Ren, W. Zhang, Y. Yu and J. Wang, “Activation Maximization Generative Adversarial Nets,” 3 2017.

[34] H. Ruiz, T. A. Etchells, I. H. Jarman, J. D. Martín and P. J. G. Lisboa, “A principled approach to network-based classification and data representation,” *Neurocomputing,* vol. 112, p. 79–91, 7 2013.

[35] I. Misra and L. van der Maaten, “Self-Supervised Learning of Pretext-Invariant Representations,” 12 2019.

[36] SimplicityBio, [Online]. Available: https://www.quartz.bio.

[37] K. Sokol and P. Flach, “Explainability fact sheets: A framework for systematic assessment of explainable approaches,” *FAT* 2020 - Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency,* p. 56–67, 2020.

[38] K. Bonawitz et al., “Towards federated learning at scale: System design,” *arXiv preprint arXiv:1902.01046.*

[39] S. Dash, S. K. Shakyawar, M. Sharma and S. Kaushik, “Big data in healthcare: management, analysis and future prospects,” *Journal of Big Data,* vol. 6, no. 54, p. 25, 2019.

[40] J. Xu and F. Wang, “Federated Learning for Healthcare Informatics,” *arXiv preprint arXiv:1911.06270,* p. 25, 2019.

[41] S. Caldas, S. M. K. Duddu, P. Wu, T. Li, J. Konečný, H. B. McMahan, V. Smith and A. Talwalkar, “Leaf: A benchmark for federated settings,” *arXiv preprint arXiv:1812.01097,* 2018.

[42] H. B. McMahan, E. Moore, D. Ramage, S. Hampson and B. A. Arcas, “Communication-efficient learning of deep networks from decentralized data,” *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS) ,* vol. 54, pp. 1273-1282, 2017.

[43] Y. Zhao, M. Li, L. Lai, N. Suda, D. Civin and V. Chandra, “Federated learning with non-iid data,” *arXiv preprint arXiv:1806.00582,* p. 13, 2018.

[44] M. Mohri, G. Sivek and A. T. Suresh, "Agnostic federated learning," *Proceedings of the 36th International Conference on Machine Learning,* vol. 97, pp. 4615-4625, 2019.

[45] L. Corinzia and J. M. Buhmann, "Variational federated multi-task learning," *arXiv preprint arXiv:1906.06268,* p. 10, 2019.

[46] R. Shokri and V. Shmatikov, "Privacy-preserving deep learning," *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security,* p. 1310–1321, 2015.

[47] T. Nishio and R. Yonetani, "Client selection for federated learning with heterogeneous resources in mobile edge," *ICC 2019-2019 IEEE International Conference on Communications,* pp. 1-7, 2019.

[48] H. Zhu and Y. Jin, "Multi-objective evolutionary federated learning," *IEEE transactions on neural networks and learning systems,* vol. 31, no. 4, pp. 1310-1322, 2020.

[49] M. Kamp, L. Adilova, J. Sicking, F. Hüger, P. Schlicht, T. Wirtz and S. Wrobel, "Efficient decentralized deep learning by dynamic model averaging," *Joint European Conference on Machine Learning and Knowledge Discovery in Databases,* vol. LNCS 11051, pp. 393-409, 2018.

[50] A. G. Roy, S. Siddiqui, S. Pölsterl, N. Navab and C. Wachinger, "Braintorrent: A peer-to-peer environment for decentralized federated learning," *arXiv preprint arXiv:1905.06731,* p. 9, 2019.

[51] Q. Yang, Y. Liu, T. Chen and Y. Tong, "Federated machine learning: Concept and applications," *ACM Transactions on Intelligent Systems and Technology,* vol. 10, no. 2, pp. 1-19, 2019.

[52] A. N. Bhagoji, S. Chakraborty, P. Mittal and S. Calo, "Analyzing federated learning through an adversarial lens," *Proceedings 36th International Conference on Machine Learning,* vol. 97, pp. 634-643, 2019.

[53] M. Abadi et al., "Deep learning with differential privacy," *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security,* 2016.

[54] S. Truex, N. Baracaldo, A. Anwar, T. Steinke, H. Ludwig, R. Zhang and Y. Zhou, "A hybrid approach to privacy-preserving federated learning," *Proceedings of the 12th ACM Workshop on Artificial Intelligence and Security,* 2019.

[55] T. S. Brisimia, R. Chen, T. Mela, A. Olshevsky, I. C. Paschalidis and W. Shi, "Federated learning of predictive models from federated electronic health records," *International journal of medical informatics,* vol. 112, pp. 59-67, 2018.

[56] L. Huang, A. L. Shea, H. Qian, A. Masurkar, H. Deng and D. Liu, "Patient clustering improves efficiency of federated machine learning to predict mortality and hospital stay time using distributed electronic medical records," *Journal of Biomedical Informatics,* vol. 99, p. 103291, 2019.

[57] W. Li, F. Milletar, D. Xu, N. Rieke, J. Hancox, W. Zhu, M. Baust, Y. Cheng, S. Ourselin, M. J. Cardoso and A. Feng, "Privacy-preserving Federated Brain Tumour Segmentation," *International Workshop on Machine Learning in Medical Imaging,* vol. LNCS 11861, pp. 133-141, 2019.

[58] [Online]. Available: https://github.com/tensorflow/tensorflow.

[59] [Online]. Available: https://stackshare.io/stackups/pysyft-vs-tensorflow.

[60] [Online]. Available: https://www.openmined.org/.

[61] T. Ryffel et al., "A generic framework for privacy preserving deep learning," *arXiv preprint arXiv:1811.04017,* 2018.

[62] [Online]. Available: https://github.com/SubstraFoundation/welcome.

[63] [Online]. Available: https://github.com/FederatedAI/FATE.

[64] K. Savić et al., "Feature selection based on community detection in feature correlation networks," *Computing,* vol. 101, no. 10, pp. 1513-1538, 2019.

[65] A. Vizitiu et al., "Applying Deep Neural Networks over Homomorphic Encrypted Medical Data," *Computational and Mathematical Methods in Medicine,* vol. 2020, p. 3910250, 2020.

[66] C. Gentry and S. Halevi, "Implementing gentry's fully-homomorphic encryption scheme," *Advances in Cryptology – EUROCRYPT 2011,* vol. LNCS 6632, pp. 129-148, 2011.

[67] N. Dowlin, R. Gilad-Bachrach, K. Laine, K. Lauter, M. Naehrig and J. Wernsing, "Cryptonets: Applying neural networks to encrypted data with high throughput and accuracy," *Proceedings of the 33rd International Conference on Machine Learning,* vol. 48, 2016.

[68] E. Hesamifard, H. Takabi and M. Ghasemi, "Cryptodl: Deep neural networks over encrypted data," *ArXiv, vol. abs/1711.05189,* p. 21, 2017.

[69] C. Orlandi, A. Piva and M. Barni, "Oblivious neural network computing via homomorphic encryption," *EURASIP Journal on Information Security,* vol. 2007, pp. 1-11, 2007.

[70] J. W. Bos, K. Lauter, J. Loftus and M. Naehrig, "Improved security for a ring-based fully homomorphic encryption scheme," *Cryptography and Coding (IMACC 2013),* vol. 8308, p. 75, 2013.

[71] M. Chase, R. Gilad-Bachrach, K. Laine, K. Lauter and P. Rindal, "Private collaborative neural network learning," *IACR Cryptology ePrint Arch.,* no. 2017, p. 762, 2017.

[72] M. Barni, C. Orlandi and A. Piva, "A privacy-preserving protocol for neural-network-based computation," *MM&Sec,* pp. 146-151, 2006.

[73] C. Juvekar, V. Vaikuntanathan and A. P. Chandrakasan, "Gazelle: A low latency framework for secure neural network inference," *Proceedings of the 27th USENIX Conference on Security Symposium,* pp. 1651-1668, 2018.

[74] A. El-Yahyaoui and M. D. Elkettani, "Fully homomorphic encryption: state of art and comparison," *International Journal of Computer Science and Information Security,* vol. 14, no. 4, pp. 159-167, 2016.

[75] S. S. Sathya, P. Vepakomma, R. Raskar, R. Ramachandra and S. Bhattacharya, "A review of homomorphic encryption libraries for secure computation," *CoRR,* vol. abs/1812.02428, p. 12, 2018.

[76] J. Cheon et al., "Homomorphic encryption for arithmetic of approximate numbers," *ASIACRYPT,* 2017.

[77] J. Fan et al., "Somewhat practical fully homomorphic encryption," *IACR Cryptology,* vol. 2012, p. 144, 2012.

[78] S. Halevi et al., "Algorithms in helib," *IACR Cryptology,* vol. 2014, p. 106, 2014.

[79] Z. Brakerski et al., "(leveled) fully homomorphic encryption without bootstrapping," *ITCS,* 2012.

[80] J. Mancuso, "Privacy-preserving machine learning 2018: A year in review," *https://medium.com/dropoutlabs/privacy-preserving-machine-learning-2018-ayear-in-review-b6345a95ae0f.*

[81] R. Popa, "Building practical systems that compute on encrypted data".

[82] R. Paillier, "Public-key cryptosystems based on composite degree residuosity classes," *EUROCRYPT,* 1999.

[83] T. Elgamal, "A public key cryptosystem and a signature scheme based on discrete logarithms," *IEEE Transactions on Information Theory,* vol. 31, p. 469–472, 1985.

[84] S. Goldwasser et al., "Probabilistic encryption and how to play mental poker keeping secret all partial information," *STOC,* 1982.

[85] C. Bosch et al., "A survey of provably secure searchable encryption," *ACM Comput. Surv.,* vol. 47, p. 18:1–18:51, 2014.

[86] R. Agrawal et al., "Order-preserving encryption for numeric data," *SIGMOD Conference,* 2004.

[87] P. Parmar et al., "Survey of various homomorphic encryption algorithms and schemes," 2014.

[88] H. Chung et al., "Encoding rational numbers for fhe-based applications," *IACR Cryptology,* vol. 2016, p. 344, 2016.

[89] B. Tsaban et al., "Cryptanalysis of the more symmetric key fully homomorphic encryption scheme," *J. Mathematical Cryptology,* vol. 9, pp. 75-78, 2014.

[90] A. Kipnis et al., "Efficient methods for practical fully homomorphic symmetric-key encrypton, randomization and verification," *IACR Cryptology ePrint Archive,* vol. vol. 2012, p. p. 637, 2012.

[91] P. M. Schwartz and D. J. Solove, "The PII Problem: Privacy and a New Concept of Personally Identifiable Information," *New York University Law Review,* vol. 86, p. 1841, 2011.

[92] L. Sweeney, "Achieving k-anonymity privacy protection using generalization and suppression," *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems,* vol. 10, no. 5, p. 571–588, 2002.

[93] D. Barth-Jones, "The 'Re-Identification' of Governor William Weld's Medical Information: A Critical Re-Examination of Health Data Identification Risks and Privacy Protections, Then and Now," *Then and Now,* 2012.

[94] J. Domingo-Ferrer and V. Torra, "A Critique of k-Anonymity and Some of Its Enhancements," in *Third International Conference on Availability, Reliability and Security*, Barcelona, 2008.

[95] C. Dwork, "An Ad Omnia Approach to Defining and Achieving Private Data Analysis," *Privacy, Security, and Trust in KDD""PinKDD 2007,* pp. 1-13, 2007.

[96] C. Dwork, "Differential privacy," in *ICALP'06: Proceedings of the 33rd international conference on Automata, Languages and Programming*, Venice, 2006.

[97] C. Dwork, F. McSherry, K. Nissim and A. Smith, "Calibrating Noise to Sensitivity in Private Data Analysis," *Theory of Cryptography,* pp. pp 265-284, 2006.

[98] r. McSherry, "Privacy integrated queries: an extensible platform for privacy-preserving data analysis," *Commun. ACM,* vol. 53, no. 9, p. 89–97, 2010.

[99] C. Dwork and A. Roth, "The Algorithmic Foundations of Differential Privacy," *Foundations and Trends in Theoretical Computer Science,* vol. 9, no. 3-4, p. 211–407, 2014.

[100] T. Zhu, G. Li, W. Zhou and P. S. Yu, Differential Privacy and Applications, Berlin: Springer International Publishing, 2017.

[101] M. Fredrikson, S. Jha and T. Ristenpart, "Model Inversion Attacks That Exploit Confidence Information and Basic Countermeasures," in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, Denver, 2015.

[102] F. D. McSherry, "Privacy integrated queries: an extensible platform for privacy-preserving data analysis," in *SIGMOD '09: Proceedings of the 2009 ACM SIGMOD International Conference on Management of data*, Providence, Rhode Island, USA, 2009.

[103] A. Friedman, "Data mining with differential privacy," in *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Washington, DC, USA, 2010.

[104] F. K. D. Emam and K. El, "Practicing Differential Privacy in Health Care: A Review," *Trans. Data Privacy,* vol. 6, no. 1, p. 35–67, 2013.

[105] B. K. Beaulieu-Jones, Z. S. Wu, C. Williams, R. Lee, S. P. Bhavnani, J. B. Byrd and C. S. Greene, "Privacy-Preserving Generative Deep Neural Networks Support Clinical Data Sharing," *Circulation: Cardiovascular Quality and Outcomes,* vol. 12, no. 7, 2019.

[106] S. Bradner, "Key words for use in RFCs to Indicate Requirement Levels," Available at: https://www.ietf.org/rfc/rfc2119.txt, 1997.

[107] European Commission, "EU data protection rules," *https://ec.europa.eu/info/law/law-topic/data-protection/eu-data-protection-rules_en,* 2018-05-25.

[108] "ISO/IEC 25010:2011," ISO, 2017. [Online]. Available: https://www.iso.org/standard/35733.html. [Accessed 27 April 2020].

[109] European Union, "Regulation 2016/679 of the European parliament and the Council of the European Union," *Official Journal of the European Communities,* vol. 2014, p. 1–88, 2016.

[110] G. Bologna and Y. Hayashi, "Characterization of symbolic rules embedded in deep DIMLP networks: A challenge to transparency of deep learning," *Journal of Artificial Intelligence and Soft Computing Research,* vol. 7, p. 265–286, 2017.

[111] H. Zhu et al., "Multi-objective evolutionary federated learning," *IEEE transactions on neural networks and learning systems,* 2019.

[112] F. Mols, M. Thong, P. Vissers, T. Nijsten and L. v. d. Poll-Franse, "Socio-economic implications of cancer survivorship: results from the PROFILES registry," *Eur J Cancer,* vol. vol. 48(13), p. 2037–2042, 2012.