Big Data technologies and extreme-scale analytics

MARVEL

Multimodal Extreme Scale Data Analytics for Smart Cities Environments

D4.3: MARVEL's decision-making toolkit – initial version[†]

Abstract: The purpose of this deliverable is to describe all activities related to the design, implementation, and release of MARVEL's Decision-Making Toolkit. This document sets the scope and the goals of the Decision-Making Toolkit, as well as discusses the process to achieve these, the using user design process, the definition of specific user journeys, and demonstration of the use cases. More specifically, the document details the user-centric design process followed for MARVEL, the internal architecture and its features, the requirements and specifications of the Toolkit, and an extensive demonstration of the selected use cases from the user point of view. Finally, a detailed description of the conclusions and the contributions to MARVEL's objectives are setting the path to the final version of the Decision-Making Toolkit.

Contractual Date of Delivery	30/06/2022
Actual Date of Delivery	30/06/2022
Deliverable Security Class	Public
Editor	*Stella Markopoulou, Dora Kallipolitou, and Spiros Fotis (ZELUS)*
Contributors	ZELUS, FORTH, IFAG, AU, CNR, FBK, AUD, TAU, STS, UNS, ITML, GRN, MT
Quality Assurance	*Bahaeddine Abrougui (AUD)* *Borja Sáez (IFAG)*

[†] *The research leading to these results has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 957337.*

The *MARVEL* Consortium

Part. No.	Participant organisation name	Participant Short Name	Role	Country
1	FOUNDATION FOR RESEARCH AND TECHNOLOGY HELLAS	FORTH	Coordinator	EL
2	INFINEON TECHNOLOGIES AG	IFAG	Principal Contractor	DE
3	AARHUS UNIVERSITET	AU	Principal Contractor	DK
4	ATOS SPAIN SA	ATOS	Principal Contractor	ES
5	CONSIGLIO NAZIONALE DELLE RICERCHE	CNR	Principal Contractor	IT
6	INTRASOFT INTERNATIONAL S.A.	INTRA	Principal Contractor	LU
7	FONDAZIONE BRUNO KESSLER	FBK	Principal Contractor	IT
8	AUDEERING GMBH	AUD	Principal Contractor	DE
9	TAMPERE UNIVERSITY	TAU	Principal Contractor	FI
10	PRIVANOVA SAS	PN	Principal Contractor	FR
11	SPHYNX TECHNOLOGY SOLUTIONS AG	STS	Principal Contractor	CH
12	COMUNE DI TRENTO	MT	Principal Contractor	IT
13	UNIVERZITET U NOVOM SADU FAKULTET TEHNICKIH NAUKA	UNS	Principal Contractor	RS
14	INFORMATION TECHNOLOGY FOR MARKET LEADERSHIP	ITML	Principal Contractor	EL
15	GREENROADS LIMITED	GRN	Principal Contractor	MT
16	ZELUS IKE	ZELUS	Principal Contractor	EL
17	INSTYTUT CHEMII BIOORGANICZNEJ POLSKIEJ AKADEMII NAUK	PSNC	Principal Contractor	PL

Document Revisions & Quality Assurance

Internal Reviewers created

1. *Borja Sáez, IFAG*
2. *Bahaeddine Abrougui, AUD*

Revisions

Version	Date	By	Overview
0.9	29/06/2022	Editors	Addressed comments from the PC
0.7	20/06/2022	Stella Markopoulou (ZELUS)	Revision after internal review and final draft for PC
0.6	10/06/2022	Bahaeddine Abrougui (AUD)	Internal review
0.5	09/06/2022	Borja Sáez (IFAG)	Internal review
0.4	06/06/2022	Editors & Contributors	Final draft for internal review
0.3	27/04/2022	Stella Markopoulou (ZELUS)	First draft for contributions by partners
0.2	18/04/2022	Dragana Bajovic (UNS), Antonio Escobar (IFAG)	Comments on the ToC
0.1	31/03/2022	Stella Markopoulou (ZELUS)	ToC draft

Disclaimer

The work described in this document has been conducted within the MARVEL project. This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 957337. This document does not reflect the opinion of the European Union, and the European Union is not responsible for any use that might be made of the information contained therein.

Table of Contents

List of Tables

List of Figures

List of Abbreviations

API	Application Programming Interface
AT	Audio Tagging
AV	Audio-Visual
AVCC	Audio Visual Crowd Counting
DFB	Data Fusion Bus
DMT	Decision-Making Toolkit
E2F2C	Edge-to-Fog-to-Cloud
FFMPEG	Fast Forward MPEG
GUI	Graphical User Interface
JSON	JavaScript Object Notation
KPI	Key Performance Indicator
MQTT	Message Queuing Telemetry Transport
MVP	Minimum Viable Product
REST	Representational State Transfer
RTSP	Real-Time Streaming Protocol
SED	Sound Event Detection
UCD	User-Centric Design
UI	User Interface
UX	User experience
VPN	Virtual Private Network
WP	Work Package

Executive Summary

This deliverable describes all activities related to the design, implementation, and release of the initial version of MARVEL's Decision-Making Toolkit (DMT), namely the interface between the end-users and the MARVEL framework. At this stage of the project, the scope of the DMT included the advanced visualisation of information that served in total five (5) use cases in the smart cities of Malta, of Trento in Italy and of the University of Novi Sad in Serbia.

Having set the scope for this initial version of the MARVEL DMT, the goal was to develop and demonstrate advanced visualisations of real-time and historical data which would allow users to:

- timely intervene in order to address a traffic anomaly or to manage traffic flow conditions real-time;
- plan urban infrastructure developments that will facilitate and/or improve mobility based on historical data about the road usage and vehicle movement;
- timely intervene in order to address and/or prevent incidents that may occur in a crowded area;
- study the crowd movement and behaviour in order to better organize large events such as festivals or better monitor regularly crowded areas such as urban markets.

To achieve those goals, the User-Centric Design methodology was utilised for the design of the user journeys as well as of the user interfaces with the MARVEL framework in order to achieve increased usability. This document describes in further detail the methodology followed and the resulting requirements and specifications for the features of the DMT, as well as for its internal architecture. An extensive demonstration of the selected use cases from the user's point of view is also included in the document. Screenshots from the actual DMT are depicting almost every step in the user journeys in scope and are accompanied by explanations about the workflows linked to each interface component.

Finally, a detailed description of the conclusions and the contributions to the objectives of the MARVEL project is provided to the reader at the end of the document, laying the path to the final version of the Decision-Making Toolkit.

1 Introduction

This document is the third deliverable of WP4 and describes the initial version of a full-fledged Decision-Making Toolkit (DMT). In comparison to the MARVEL Minimum Viable Product (MVP) described in Deliverable 5.1[1], this version includes a user interface and visualisation functionalities for the five MARVEL use cases that were planned for M18 rather than the one use case that was included in the MVP.

1.1 Purpose of this document

The purpose of this deliverable is to describe the design rationale, the technical details, and the requirements of the initial version of the DMT in order to serve the initial version of MARVEL's Integrated framework. Within the context of MARVEL, the initial version of the Integrated framework is the release after the MVP that served as a proof-of-concept prototype for the project's main objectives.

This document stands upon the detailed description of the DMT's capabilities and needs to design the expected interactions, user journeys, requirements, and design. The work, which is presented in this deliverable, strongly relates to the user-centric design (UCD) process that the DMT follows in order to ensure that the development is aligned to the user needs and the released software receives the acceptance of the stakeholder groups.

1.2 Contributions to WP4 and project objectives

Task 4.4 and in particular the development of the Decision-Making Toolkit contributes to the fourth objective of WP4 namely, "*the development of advanced visualisation techniques to support both real-time and long-term decision-making*". These techniques include:

- Analytical reasoning for medium to long-term business decision-making based on queries execution over the processed audio-visual data;
- Data presentation and advanced visualisations that reveal hidden insights of valuable knowledge;
- Text-annotated attention maps which will enhance video streams with textual information and indications of associated audio events;
- Multisource, multimodal summaries, that allow users to explore and understand audio-visual, sensor, and other context-enriching data (e.g., weather data, information from incident reporting systems, parking sensors, etc.) and interact with them;
- Real-time visualisations of alerts and detected events for short-term decisions and monitoring, supported by a rule-based engine.

All of the above techniques have been applied to the current version of the Toolkit, with some of them, such as the analytical reasoning, the text-annotated attention maps, and the real-time visualisations of alerts, being at an advanced level, while others being at an earlier stage of development such as the hidden insights revealing visualisations and the multisource, multimodal summaries.

The work performed in this task directly addresses the project KPI-O3-E4-1: Detailed insights into more than 5 hidden correlations. This KPI is one of the six (6) KPIs quantifying the progress towards Objective 3 of the project: *"Break technological silos, converge very diverse and novel engineering paradigms and establish a distributed and secure Edge-to-Fog-to-Cloud (E2F2C) ubiquitous computing framework in the big data value chain."* It directly links with

[1] "D5.1: MARVEL Minimum Viable Product" Project MARVEL, 2021. https://doi.org/10.5281/zenodo.5833310

Enabler 4-Complex decision-making and insights, according to which the project will introduce innovative decision-making tools by combining several complementary techniques to reveal hidden relationships among the captured data.

1.3 Relation to other WPs, deliverables, and activities

The work presented in this deliverable, strongly relates to the UCD methodology that the project follows to develop the MARVEL functionalities, aiming to ensure that the development is aligned to the user needs and the released software receives the acceptance of the stakeholder groups. As such, this deliverable establishes relationships with the work performed in WP6 regarding the evaluation and validation of the DMT usability and user-friendliness through real-life industrial experiments. It is also strongly connected with technical WPs (WP2, WP3, and WP5) since the described utilised visualisations and multipurpose interfaces receive the outputs of the tools developed under these WPs and integrated into the MARVEL framework.

1.4 Structure of the document

The document consists of 7 sections as follows:

- Section 1 provides an overview of the deliverable and how it is linked to the overall project as well as to the rest of the Work Packages.
- Section 2 describes the UCD process utilised in order to achieve a high degree of usability for the Decision-Making Toolkit. A key requirement for this was to ensure the engagement of the MARVEL end-users in the implementation as early as possible in the project time plan. This document presents the outcome of this action.
- Section 3 presents the main components used to create the DMT as well as its internal architecture.
- Section 4 gives an overview of the user journeys supported by the current version of the DMT, in order to facilitate the realisation of the defined pilot use cases.
- Section 5 includes screenshots from the initial version of the DMT, following the multiple user journeys. This version will also be used for the second round of the DMT evaluation by the end-users.
- Section 6 revisits the project-related KPIs that Task 4.4 contributes to, as well as the specific component KPIs and showcases their progress.
- Section 7 presents the conclusions of the work performed and the future steps, and wraps up the document.

2 User-Centric Design

2.1 User-centric design process

The development of a software system follows the software life cycle processes defined in the ISO/IEC 12207:2017 standard[2]. This standard divide software implementation into many phases, including software requirement analysis, software architectural design, software construction and integration, software qualification testing, and software support and reuse. The support services include the software quality assurance procedures and validation processes. Software validation is the analysis and providing of objective proof that the software specifications adhere to user demands and intended purposes, and that the specific requirements implemented through software can be consistently met. Since the software is usually part of a larger hardware system, the software validation typically includes evidence that all software requirements have been implemented correctly and completely.

The degree to which stakeholders are involved in the processes of a software implementation project determines whether a UCD approach is used in the creation of the software system. A UCD methodology, also known as human-centred design, aims to bring together a multidisciplinary team to work on the interactive creation of software solutions with a focus on usability validation. As stated in ISO 9241-210:2010[3], UCD is the approach to systems design and development that aims to make interactive systems more usable by focusing on the use of the system and applying human factors/ergonomics and usability knowledge and techniques. Consequently, the scope of a UCD process is to achieve a high degree of usability, which is introduced as an inherent measurable property for all interactive digital technologies and is defined in ISO 9241-11:2018[4] as "the extent to which a system, product or service can be used by specified users to achieve specified goals with effectiveness, efficiency, and satisfaction in a specified context of use". To achieve this goal, the UCD process requires that software end-users and solution domain stakeholders be involved in the implementation activities early on, through a structured and systematic approach to gathering all the necessary information, which will guide the software designers in making an effective first approach of the software sketch (Jeffrey Rubin et. Al, 2008).

Since the UCD method is heavily reliant on user acceptance criteria, it should consider usability from a variety of angles, focusing on issues such as simplicity of use and learning curve to integrate the solution into current business processes[5]. This may be assessed by the intended users through an empirical study of similar solutions and real testing of the software product as it is gradually made accessible, starting with a mock-up version, and progressing through functional prototyping and product testing. In this context, iterative design is a key function in the whole software development process for the continuous evaluation of usability aspects of the software solution.

The purpose of a usability evaluation is to determine how easy and comfortable it is to use an interactive software solution. There are several approaches for determining the degree to which a software interactive solution has been developed. These approaches provide the theoretical

[2] International Organization for Standardization / International Electrotechnical Commission (ISO/IEC) 12207:2017, "Systems and software engineering - Software life cycle processes", November 2017. https://www.iso.org/standard/63712.html

[3] International Organization for Standardization (ISO) 9241-210:2010, "Ergonomics of human-system interaction -- Part 210: Human-centred design for interactive systems", March 2010. https://www.iso.org/standard/52075.html

[4] International Organization for Standardization (ISO) 9241-11:2018, "Ergonomics of human-system interaction - Part 11: Usability: Definitions and concepts", March 2018. https://www.iso.org/obp/ui/#iso:std:iso:9241:-11:ed-2:v1:en

[5] Notes on User Centred Design Process, https://www.w3.org/WAI/redesign/ucd

foundation for the use of certain robust, objective, and trustworthy measurements that enable a quantified approach to usability. The most well-known usability evaluation models are: i) the quality in use model, which is used to assess the use of the software product (effectiveness, efficiency, and satisfaction in a particular context of use), and ii) the product quality model, which is used to assess the product user interface and interaction. The ISO/IEC 12207:2017 standard[6] includes these quality models as part of the software validation process. However, in a UCD process, these models should be used throughout the software implementation phase.

The ISO/IEC 25010:2011 standard[7] is the most widely used reference model for software assurance. It provides eight software quality criteria that may be used to evaluate a software solution. The aforementioned quality models are introduced in this standard, which provides a unified vocabulary for describing, measuring, and assessing system and software product quality. Thus, in ISO/IEC 25010:2011 standard:

- The quality in use model is composed of five characteristics that relate to the outcome of the interaction with the system and characterises the impact that the product can have on the stakeholders. In this model, usability is at the forefront and is realised through the characteristics shown in Figure 1.
- The product quality model is composed of eight characteristics that relate to static properties of software and dynamic properties of the computer system. In this model, the eight characteristics can be further divided into sub-characteristics, as shown in Figure 2.

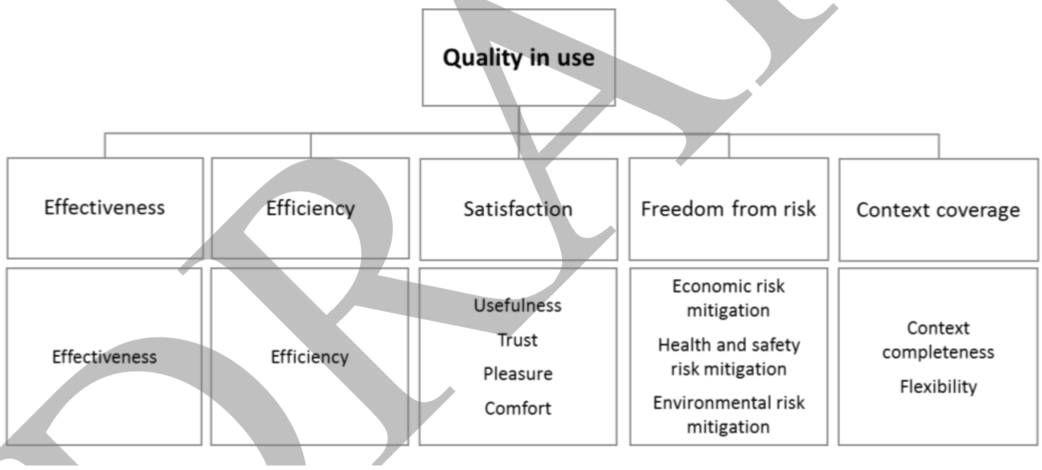

Figure 1: The quality in use model of the ISO/IEV 25010:2011 standard

[6] International Organization for Standardization / International Electrotechnical Commission (ISO/IEC) 12207:2017, "Systems and software engineering - Software life cycle processes", November 2017. https://www.iso.org/standard/63712.html

[7] International Organization for Standardization / International Electrotechnical Commission (ISO/IEC) 25010:2011, "Systems and software engineering - Systems and software Quality Requirements and Evaluation (SQuaRE) - System and software quality models", March 2011. https://www.iso.org/standard/35733.html

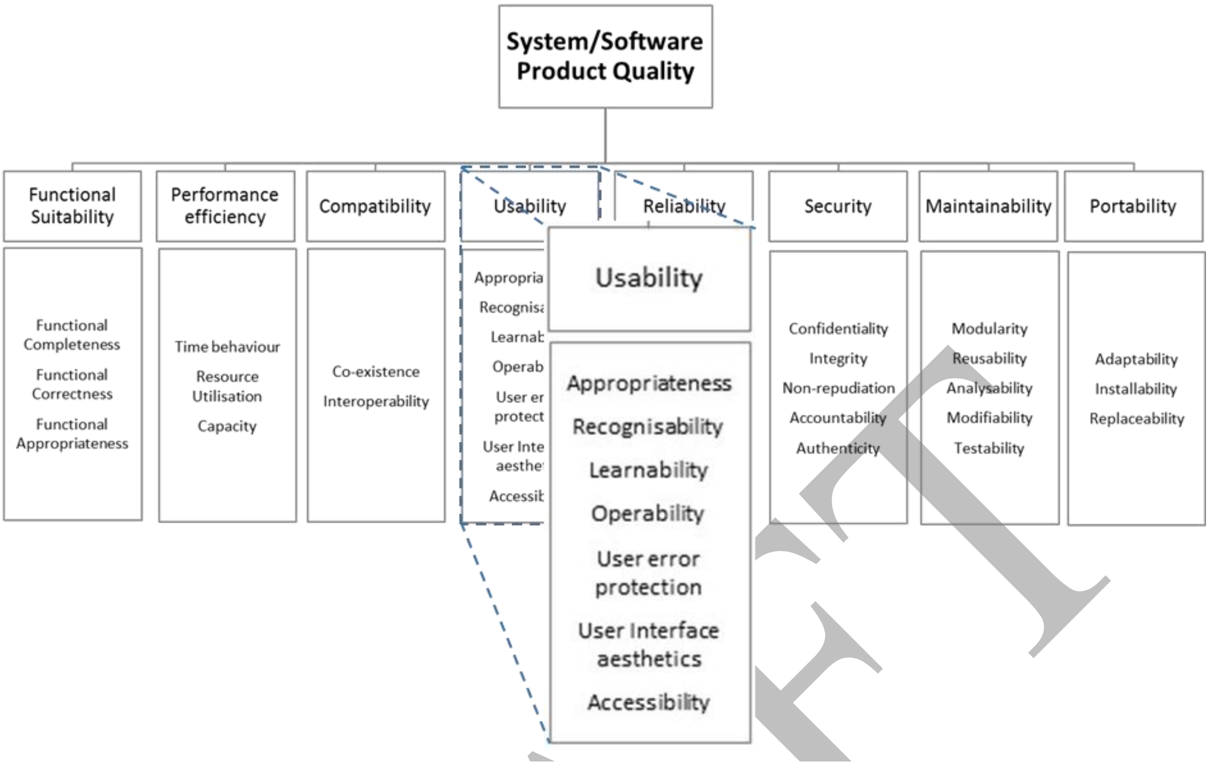

Figure 2: The product quality model of the ISO/IEC 25010:2011 standard

The software quality model of the ISO/IEC 25010:2011 standard is adapted on a case-by-case basis, in order to define appropriate metrics and to be able to evaluate the software capabilities. These measures must accurately reflect the attributes they represent. They must also enable relevant measures to be collected, either quantitatively (through software testing/simulations, usability tests) or qualitatively (e.g., through user observations).

Three types of classes of metrics are defined in this standard:

- Internal metrics associated with static internal properties of a system such as several function calls, and several rules.
- External metrics associated with dynamic external properties. These are metrics that are observable when the user interacts with the system (i.e., the user performs a task/function/operation and observes the response in the sense of time required, results obtained etc.).
- Quality-in-use metrics, refer to metrics that evaluate the extent to which a system meets the needs of the user.

Various standardisation groups have created common formats for usability specialists, end-users, and software development teams to produce usability requirements, focusing on usability criteria. The ISO 25062:2006 standard for Uniform Industry Format for Usability Test Reports[8]

[8] International Organization for Standardization (ISO) 25062:2006: "Software engineering - Software Product Quality Requirements and Evaluation (SQuaRE) - Common Industry Format (CIF) for usability test reports", April 2006. https://www.iso.org/obp/ui/#iso:std:iso-iec:25062:ed-1:v2:en.

and the NISTIR 7432: Common Industry Specification for Usability – Requirements (CIRU-R)[9] both provide such common forms.

2.2 User-centric design process for MARVEL

This section outlines MARVEL's approach to creating the interfaces that will be utilised by the intended stakeholders, considering the UCD process's core phases as previously outlined.

The analysis of usability aspects, the design of interface components with usability criteria as part of the overall MARVEL platform specifications, and ultimately the usability evaluation of the integrated environment, are all part of the user participation in the MARVEL implementation phase. Deliverable D1.3[10] has previously given the description of end-user groups, their demands, and the following functional and non-functional requirements for the definition of MARVEL functionalities as well as the context of the use of these functions by the targeted end-users.

In this deliverable, we focus on the visual representation of MARVEL concepts and the anticipated wireframes that will eventually constitute the user interface components of the MARVEL integrated platform. The wireframes depict a preliminary picture of how the defined end-users would interact with the MARVEL environment in order to manage their datasets and conduct the available analyses in order to derive meaningful insights. Therefore, considering the groups of end-users, the use cases of each pilot that are to be accomplished via MARVEL, the type of data to be collected, processed, and visualised within the MARVEL environment and the expected user requirements to be fulfilled by the MARVEL platform, we evolve the design of the user interface wireframes around the following principles:

- Appropriateness: The wireframes should represent the interfaces' appropriateness to allow interaction with the MARVEL environment by the various intended end-user groups i.e., non-IT specialists, data scientists, etc.
- Operability: The wireframes should enable end-users to participate autonomously by providing self-explained interactive tools with simple operation and control.
- Flexibility: Through extensive navigation, the wireframes should allow end-users to take control of the workflow execution.
- Awareness: The wireframes should highlight the MARVEL platform's capacity to show to the end-users the progress of the process and the consistency of the information provided through the interface functions.
- Efficiency: the wireframes should provide clarity on which is the primary action of each wireframe that should be accomplished by the end-user.
- User interface (UI) aesthetics: the wireframes should be customised to the behaviour and philosophy of the end-users using modern design elements so that they are visually pleasant and attractive for use.
- Accessibility: UI development will be based on international standards and will target various access channels (end-user devices) as well as for as many users as possible.

It should be mentioned that relevant work in similar systems that handle Big Data analytics will be considered while creating the MARVEL User Interface. However, MARVEL seeks to provide a sophisticated visualisation framework based on new revolutionary technologies and existing commercial frameworks, like ZELUS' SmartViz. Furthermore, MARVEL's UI features

[9] National Institute of Standards and Technology (NIST), NISTIR 7432: "Common Industry Specification for Usability – Requirements (CIRU-R)", June 2007. https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=51179

[10] "D1.3: Architecture definition for MARVEL framework," Project MARVEL, 2020. https://doi.org/10.5281/zenodo.5463897

differ from other largely IT-expert oriented platforms in that they attempt to satisfy the demands of both non-IT and IT professionals using the same platform. Visualisation options utilised in similar tools/platforms, on the other hand, will serve as beneficial inspiration during talks with users about what is the best match for their needs.

3 Decision-Making toolkit's Architecture

3.1 Decision-Making toolkit

The DMT is the key means of interaction with MARVEL end-users. It is serving as the key interface of all processes performed in MARVEL so that the users can consume the extracted insights. Receiving input from the Data management and distribution subsystem, as well as insights from Edge-Fog-Cloud (E2F2C) deployed models that will allow the visualisation of recommendations for medium to long-term decisions, the DMT can be envisioned as a dashboard with pre-configured widgets according to the preferences of users.

SmartViz is a data visualisation toolkit that constitutes the UI of the DMT. It provides the means to visualise several incoming detected events and anomalies deriving from the analysis of incoming data through the edge layer. It consists of a set of visualisation tools developed to allow an exploratory analysis of the incoming data by using interactive representations of data. SmartViz enables the end-users to interact with and gain a solid understanding of data, to drill into more detailed information, to discover patterns and correlations of data items and to lead them in making data-driven decisions. SmartViz functionalities include advanced visualisations of detected events, demonstration of the related statistical data via different visualisation widgets, data filtering options, and flexible interfaces.

SmartViz is also able to consume a variety of data coming through different workflows inside the MARVEL infrastructure. Such data has different velocities and formats and can be offered via live streams or a historical database, allowing the MARVEL DMT to be able to handle and support all the above and in parallel to satisfy the end-user needs and facilitate their decision-making.

3.2 Technologies and Architecture

SmartViz can visualise either historical or real-time data. Based on its modular internal architecture, presented in Figure 3: Decision-Making Toolkit architecture, it can be used to provide visualisation configurations, covering the needs of a variety of user paths that might be required in order to address all MARVEL stakeholders. Using its Data Intake adapters, SmartViz is capable of connecting with multiple data sources and then using its internal data API and configuration options to produce predefined as well as user-defined visualisation dashboards. Existing adapters can already support RESTful APIs and real-time data feeds (i.e., Kafka topics, MQTT). Furthermore, more data adapters can be plugged in to support different data sources if needed. The output of the adapters is handled by a middleware, that transforms information into internal data representations, which can afterwards feed the visualisations. This middleware also hosts configuration options, managing parameters like user authentication and authorisation, dashboard sharing options, interface layout options, etc.

The Frontend part of the tool is served as a web application directly accessible by end-users. Taking advantage of the described setup, data can be displayed using a pool of visualisations. Such visualisations can include elements varying from simple charts and graphs to more complex timeline representations, geospatial depictions, and real-time rendering of data streams. Moreover, advanced filtering mechanisms, interconnected visualisations, and data comparison features are available to allow multiple ways of data presentation and foster exploratory data analysis.

Using its capabilities, SmartViz aims to help users greatly decrease the time needed to gain a solid situational awareness. Therefore, decision-makers will be empowered to discover patterns, behaviours, and correlations of data items via a visual data exploration process that will be fed by the available information collected and/or processed by the rest of the MARVEL subsystems.

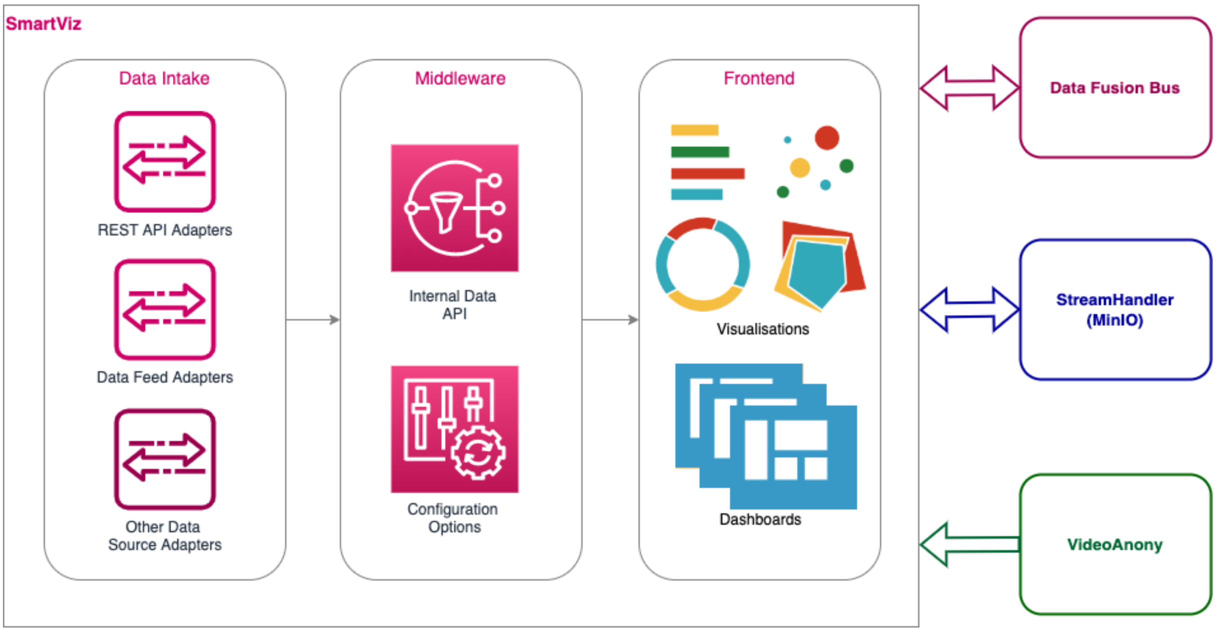

Figure 3: Decision-Making Toolkit architecture

3.3 Features

SmartViz carries out the visualisations through a pool of visualisation widgets. The widgets used for the MARVEL's DMT are carefully selected in order to satisfy both users and data needs. The widgets that are assigned to each use cases are displayed in Table 1 below.

Table 1: Assigned widgets per use case

Widgets	GRN3	GRN4	MT1	MT3	UNS1
Temporal Representation	-	YES	YES	YES	-
Crowd Density Heatmap	-	YES	-	-	YES
Details Widget	YES	YES	YES	YES	-
Vehicle Trajectories	-	YES	-	-	-
Statistical Representation	YES	YES	YES	-	-
Video Stream Player	YES	-	YES	YES	YES
Real-time Map Representation	YES	-	YES	-	-
Summaries	-	-	-	YES	YES

In the following subsections, we describe the requirements, the functionalities, and the operations of the widgets used so far in the visualisation toolkit.

3.3.1 Temporal Representation

The Temporal Representation widget shown in Figure 4, consists of a representation of detected events and anomalies across time. This timeline consists of two axes with the horizontal one containing the timestamps divided into slots relative to the events' distribution, and the vertical one having the groups of the items. This widget can be used to represent a variety of information. In the scope of the MARVEL project and the selected use cases, this widget represents detected vehicles crossing street junctions or events and anomalies detected by the MARVEL components conducting data collection and analysis. The groups that the timeline contains are colour-coded, labelled with the type of vehicle or anomaly, and represented as dots at the exact moment of detection. The user who interacts with this widget can zoom in and out in time, pan left and right in order to see all the events contained in it, double click on the events to add some important information in the filtering options, and finally hover on the events to see additional available information. Evidently, this widget only works with data associated with a timestamp.

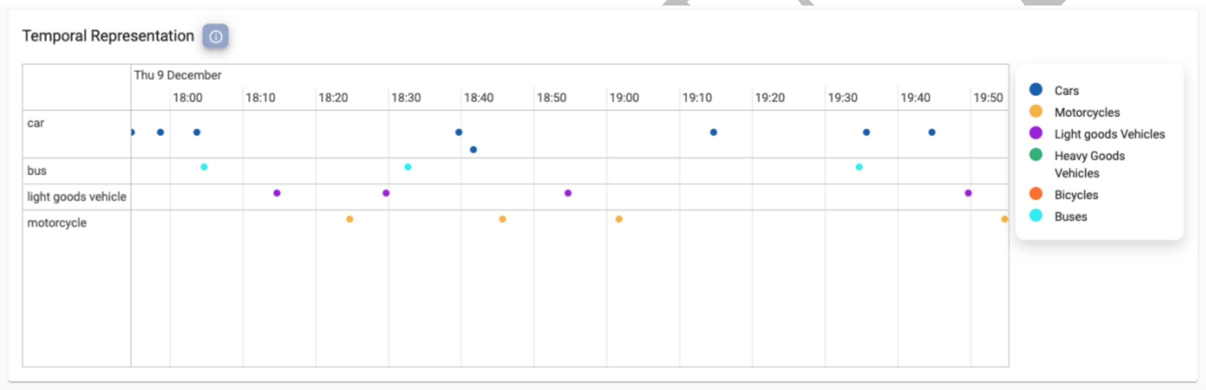

Figure 4: Temporal Representation widget

3.3.2 Crowd Density Heatmap

The Crowd Density Heatmap shown in Figure 5, is a widget developed solely for the representation of the output analysis of the Audio-Visual Crowd Counting (AVCC) and Visual Crowd Counting (VCC) components developed by AU. This feature is the visualisation of the analysis of a video frame and the crowd it contains. By constructing a heatmap on top of an original video frame, we can showcase the probability of a pedestrian being present at a moment in time, with green being the lowest and red the highest. The ability to replay the consecutive created heatmaps from a selected time period is also offered as a way for the user to see the crowd density and distribution across time. Finally, a numerical indicator of the detected pedestrians in the frames is also displayed on the right side of the widget. The frames used behind the created heatmap are static images collected by the camera frames with limited vehicles and pedestrians in them, taken at different hours of the day and night to give more context as well to the user analysing them.

Figure 5: Crowd Density Heatmap widget

3.3.3 Details Widget

The Details Widget depicted in Figure 6, presents the incoming events, anomalies, and any type of available and important information in a tabular form. This widget is used to give the user the opportunity to drill into more detailed information regarding the events that are visualised in a different way in the DMT. This widget also supports a standalone text filtering capability in order to search through the available information.

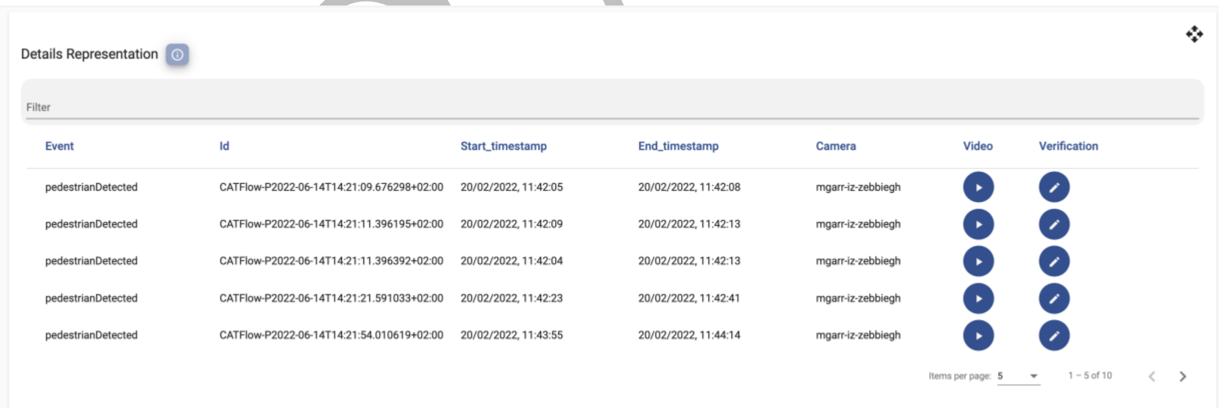

Event	Id	Start_timestamp	End_timestamp	Camera	Video	Verification
pedestrianDetected	CATFlow-P2022-06-14T14:21:09.676298+02:00	20/02/2022, 11:42:05	20/02/2022, 11:42:08	mgarr-iz-zebbiegh		
pedestrianDetected	CATFlow-P2022-06-14T14:21:11.396195+02:00	20/02/2022, 11:42:09	20/02/2022, 11:42:13	mgarr-iz-zebbiegh		
pedestrianDetected	CATFlow-P2022-06-14T14:21:11.396392+02:00	20/02/2022, 11:42:04	20/02/2022, 11:42:13	mgarr-iz-zebbiegh		
pedestrianDetected	CATFlow-P2022-06-14T14:21:21.591033+02:00	20/02/2022, 11:42:23	20/02/2022, 11:42:41	mgarr-iz-zebbiegh		
pedestrianDetected	CATFlow-P2022-06-14T14:21:54.010619+02:00	20/02/2022, 11:43:55	20/02/2022, 11:44:14	mgarr-iz-zebbiegh		

Figure 6: Details widget

3.3.4 Vehicle Trajectories

In the Vehicle Trajectories, shown in Figure 7, the paths of the passing vehicles are depicted in a frame from a traffic-monitoring camera feed recording. This widget is developed in the scope of the MARVEL project for visualising the output analysis driven by the CATflow component developed by GRN. This widget consists of the depiction of the passing vehicles and their direction, drawn in a video frame colour-coded by their type. The user is able to see an animated version of how the vehicle trajectories are building up, apply filters in order to only see specific vehicle types or directions, and finally adjust the slider on the bottom to a specific time period for further investigation.

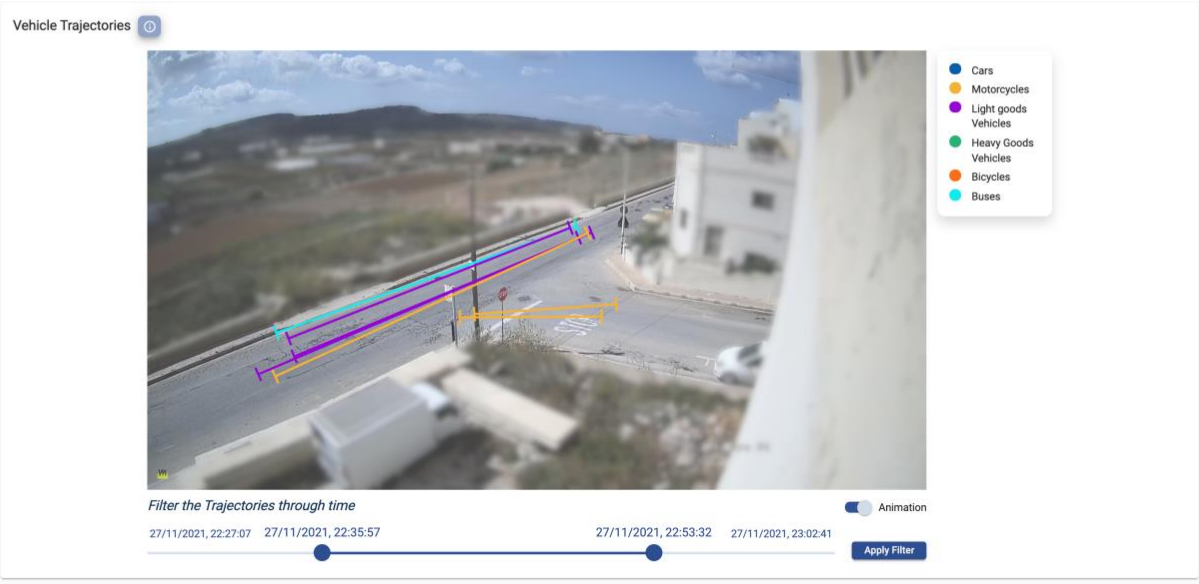

Figure 7: Vehicle Trajectories widget

3.3.5 Statistical Representation

The Statistical Representation widget, displayed in Figure 8, consists of many charts such as line, bar, and pie charts and textual information, carefully selected for each use case according to the pilots' and users' needs. This widget, similarly to all the widgets in a dashboard, is timewise synchronised and connected to the available filters offered by the DMT. In the context of MARVEL, this widget is representing different data in each use case and is mostly connected to the percentages and statistics for detected vehicles, pedestrians, events, anomalies, and additional available information for their attributes.

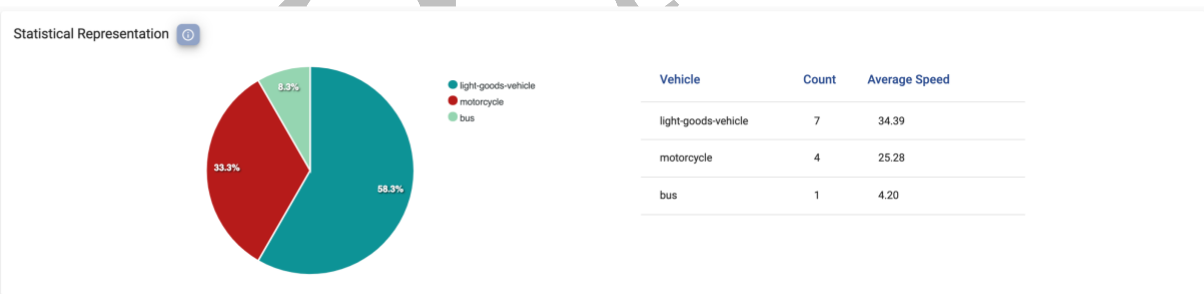

Figure 8: Statistical Representation widget

3.3.6 Video Stream Player

The Video Stream Player widget shown in Figure 9, is solely developed for the MARVEL project. This widget is expecting an RTSP URL of either a camera live stream or a video segment extracted from a camera feed that falls under the time period selected by the user. The user is able to select the time period the video has occurred, which camera feed will be displayed, and finally to play, pause, and stop the displayed video stream.

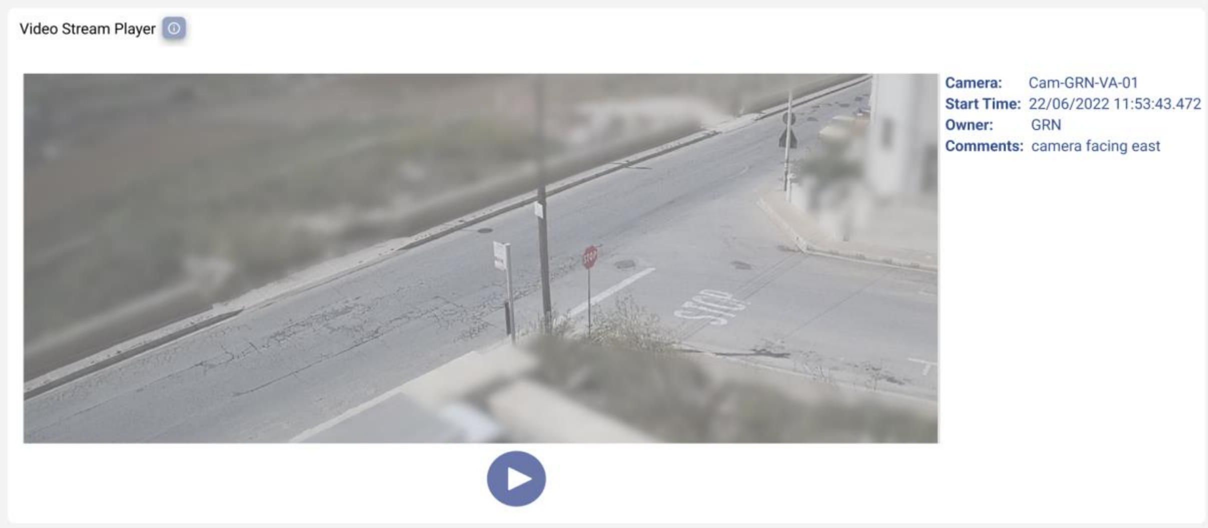

Figure 9: Video Stream Player

3.3.7 Real-time Map Representation

The Real-time Map Representation widget depicted in Figure 10, is showcasing the real-time detected anomalies and events. This widget receives event type data by a Kafka topic and presents them as pins on top of the map. The user is able to click on a pin on the map to see additional information regarding the corresponding event and the corresponding video displayed on the Video Stream Player widget described in Section 3.3.6.

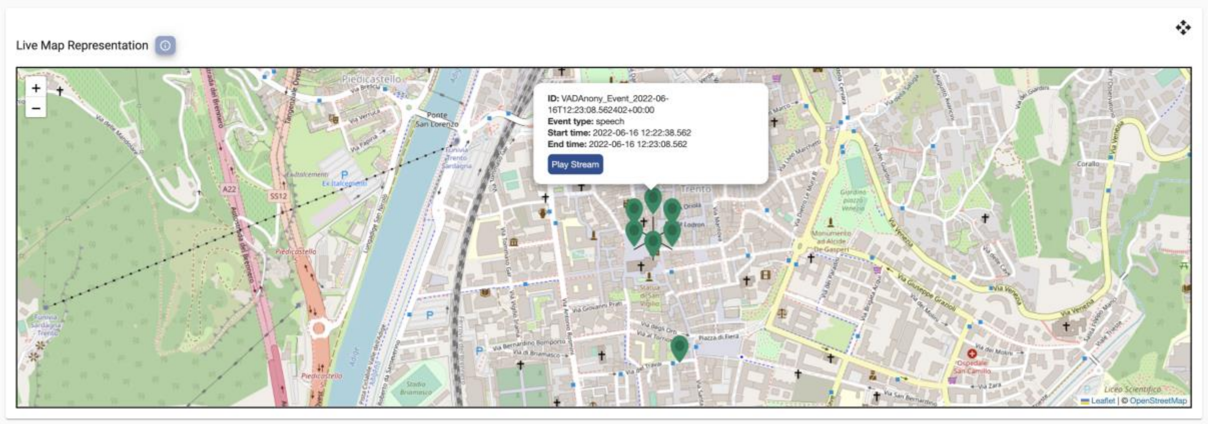

Figure 10: Real-time Map Representation widget

3.3.8 Summaries

The Summaries widget, depicted in Figure 11, is used to indicate the detected vehicles and/or pedestrians depending on the use case's needs. This widget for M18 will be used in two use cases, the MT3 – Monitoring of parking places and the UNS1 – Drone Experiment. These use cases are both linked to the number of detected vehicles and pedestrians in the monitored areas. In the future, it will be extended to include representation of summaries of pictures extracted by the FedL component to curate a story for the user.

Figure 11: Summaries widget

4 Initial requirements and specifications for the DMT

As it was already mentioned in Section 2, for the design and development of the DMT, it was made sure that the users and their needs were taken into consideration early in the development process, in order to ensure the DMT's usability, user-friendliness, and most importantly its relevance to successfully address the user's needs. Building on the MVP's functionalities, the initial version of the DMT facilitates the decision-making process in five different use cases: four related to the smart cities needs of Malta and Trento and one related to the experiment at the campus of the University of Novi Sad, as described in the following chapters.

4.1 Selected use cases

Within MARVEL 10 use cases will be realised offered by the three participating pilots. There are in total four use cases for the smart city of Malta in a pilot led by GRN. For the initial version of the DMT, use case GRN3, *"Traffic Conditions and Anomalous Events"* and use case GRN4, *"Junction Traffic Trajectory Collection"* were addressed. The latter one is an update of the use case tackled in the MVP. In the case of the smart city of Trento, there are also four use cases in the pilot led by the Municipality of Trento in close collaboration with FBK. The ones addressed in this version of DMT are use case MT1, *"Monitoring of crowded areas"* and use case MT3, *"Monitoring of parking places"*. Last but not least, there are in total two use cases to be addressed within the context of Novi Sad, and for this version, we tackled use case UNS1, *"Drone experiment"*.

4.1.1 GRN3 – Traffic Conditions and Anomalous Events

In this use case, traffic conditions are monitored to detect anomalous events. Such events include traffic jams, accidents, cars stuck and obstructing a junction, very slow vehicles and service vehicles parked on the side or obstructing a carriageway. The latter event is frequent in Malta's narrow one-way urban streets, often causing ripple effects that extend beyond the immediate area. In general, this output would find application in systems intended to inform drivers near the detected anomaly or to infer possible issues in adjacent areas, thus informing drivers of obstacles ahead. In addition, the detection of anomalous events can be used to alert personnel stationed at traffic management control rooms, who can then interpret the data and take the necessary action. The latter application is particularly attractive since it promises to reduce the detection time of anomalous events.

The main aim of this use case is to detect anomalous road conditions which may be related (passively or actively) to obstructions. In terms of evaluation metrics, accuracy and detection time are two crucial features of the use case and these can be used to benchmark the system as follows:

- Correct detection of the cause of obstructions on the road 70% of the time, as verified by manual processing.
- Detection of anomalous events within 2 minutes from their start, as verified by manual processing.

The upcoming integration of components in M18 is designed such that both these goals can be addressed. For this integration whenever an anomaly occurs, the anomaly is flagged (typically to the traffic personnel in the control room such that action can be taken). In this case, when an anomaly is detected, a video clip of the anomaly is saved and this could be made available to the traffic control personnel, thus gaining visual insight into the anomaly, which helps determine the right course of actions. This step is essential for the evaluation of both goals mentioned above since the traffic personnel would be able to flag the accuracy of the anomaly detection models. In addition, the timestamps on the video stream will indicate the latency with which the anomaly was detected.

The users of this system are intended to be traffic managers who can give directives to authorities to react to a traffic incident. The User Stories for this use case are shown in Table 2 and these give a better understanding of the functionality expected by the intended user.

Table 2: GRN3 – Traffic Conditions and Anomalous Events

Title	Main user stories	Sub user stories and user intent
Anomaly detection	**As a** Traffic Manager, **I want** to be alerted if an anomaly occurs on the road **so that** I can analyse it and determine a course of actions.	**As a traffic manager:** - once informed of anomalous events on the roads, I can take appropriate actions, i.e., dispatching the relevant civil protection authorities. - once an anomaly is flagged, I can see a live feed from the camera where the event occurred, as well as the feed a few minutes before the anomaly happened, to accurately assess the cause. - once I review the camera feed of a flagged event, I can have the option to mark this event as anomalous or not and continuously improve the system and in this case: i. allowing for the augmentation of additional annotated data which is needed for training or testing of AI models, and ii. measuring the accuracy of the models and in general the evaluation of the system. - if anomalies happened in the past or in my absence, I can still view them and access historical data around them, therefore such data are stored and managed. - since I am observing multiple locations, I can see observed anomalies and related data on a map.
Traffic Conditions	**As a** Traffic Manager **I want** to monitor the traffic conditions easily and efficiently on various roads and junctions **so that** I can monitor events such as traffic jams, congestion levels and traffic flow rate.	By monitoring the traffic conditions of various roads simultaneously, I can observe traffic jams, traffic bottlenecks, traffic flow rates, and obstructions. Here, as the traffic manager, I can be alerted whenever the traffic conditions or states change substantially. With the knowledge that a road is obstructed, I can alert the appropriate authorities. With long-term data, infrastructural changes that might be needed can be flagged, for further processing by the traffic and infrastructure engineers. If more than one location is being observed, this data can be represented on a map.

4.1.2 GRN4 – Junction Traffic Trajectory Collection

Junction Traffic Trajectory collection is focused on the requirement of long-term data analytics that shed light on both the behaviour of road users (e.g., car drivers, motorcyclists, cyclists, pedestrians, etc.) and on gathering traffic statistics at road network junctions. This use case is of interest for long-term transport planning and evaluation. In particular, there is currently significant interest in studying active travel modes, such as cycling, walking, and micro-mobility more generally. Authorities in Malta are interested in, for example, finding the optimal position of pedestrian crossings, whether provisions for cyclists at complex junctions are adequate, and whether installed provisions are being used as intended.

This use case requires entity detection and its trajectory across a junction or road segment and descriptive statistics of network junction traffic. It, therefore, follows that entity detection and tracking models can be potentially used as a first processing stage followed by further processing to generate descriptive statistics.

The innovation that is being targeted with this use case is the construction of a query-able database that can be used to look up historical data on the vehicle movements and flows, as well as pedestrians in junctions (e.g., roundabouts), with a sufficient accuracy to detect anomalous patterns autonomously 50% of the time.

This goal will be partially addressed in the R1 integrated framework. The trajectories and data generated from the CATFlow algorithm are saved on the MARVEL Data Corpus such that the data can be accessed and processed by the end-user. Currently, the anomalous paths entities take can be manually detected through visual inspection of the trajectories, which is a feature of the system. Future integrations will have tools that automatically detect anomalous paths. This use case was also partially implemented for the MVP at M12. The progress since the MVP implementation included the incorporation of all the components as part of the first integrated system.

The users for this system are intended to be traffic engineers who need data to make informed decisions about infrastructure changes and upkeep, as well as transport researchers who are interested in user behaviours. The User Stories for this use case are shown in Table 3 and give a better understanding of the functionality expected.

Table 3: GRN4 – Junction Traffic Trajectory Collection

Title	Main user stories	Sub user stories and user intent
Vehicle and Pedestrian Trajectories	**As a** Traffic Engineer, **I want to** view the trajectories of vehicles across a junction, **so that** I can analyse the most preferred paths vehicles take and make decisions on infrastructure. For example, whether large vehicles should be allowed in a particular street.	**As a traffic engineer:** - Once I view the most common trajectories taken by each type of vehicle or entity, I am able to make decisions on any infrastructure upgrade needed at each location. For example, I can note that pedestrians cross the street frequently from one spot, thus I can plan a crossroad at a certain point. In another case, I can note that cyclists take a certain path, and thus plan a cycling lane accordingly. The trajectories can also give insights into frequent recurring obstructions (i.e., cars parked at an inappropriate spot) which force temporal changes in the trajectories.
Vehicle and Pedestrian Counting	**As a** Traffic Engineer, **I want to** view the number of different types of vehicles that use certain types of roads, **so that** I am able to track the frequency of road usage per vehicle.	**As a traffic engineer:** - I want to easily view the traffic counts of all types of vehicles per vehicle, per time period and per road or junction, so that I can track the usage of a road per vehicle and estimate when and how the infrastructure needs to be upgraded. For example, stronger materials might be needed if the roads are used frequently by heavy vehicles or maybe wider pavements are needed if the road is a frequent walking route.

4.1.3 MT1 – Monitoring of crowded areas

This use case is focused on the selection of views of areas of interest for monitoring situations such as exceptional crowd, suspect or unusual crowd movements, etc. The areas for this scenario are Piazza Fiera, a square hosting the “Christmas Markets”, in Trento and Piazza Duomo, where the weekly market is located. Both locations host other crowded events. On these occasions (when the squares are crowded) the number of robberies and aggressions can increase. In addition, first aid may be needed for people who are unwell or faint. In order to be informed of these actions in time, alerts should be sent to the local police operational centre as quickly as possible. Visual analysis should be carried out in real-time and on recording data saved in the servers of the Local Police.

During the phase of training, videos will be collected from the cameras in Piazza Duomo and Piazza Fiera by FBK. The upload function is a secure transmission by VPN Access between MT and FBK in which raw video will be sent to the data lake in FBK. The videos are anonymised by FBK and annotated by MT. The data are stored in long-term cloud storage provided by PSNC inside the MARVEL Framework. In parallel, the Local Police has to store the raw video in the current surveillance system and data are being deleted after 7 days (this is by regulation on the investigation).

At this point, we have videos with annotations, and they can be used for developing an accurate model for detecting anomalous crowding situations as well as for the training. The results of the deployment will produce classifications of the selected event, insights related to the number of events, and other information requested by Local Police and prediction in order to understand if a possible crowded situation is dangerous or not.

The users for this use case are intended to be Local Police officers who need data to make informed decisions and monitor crowded areas. The User Stories for this use case are shown in Table 4 and give a better understanding of the functionality expected.

Table 4: MT1 – Monitoring of crowded areas

Title	Main user stories	Sub user stories and user intent
Exceptional crowd, suspect or unusual crowd movements	**As a** Local Police officer, **I want** to be alerted if an anomaly, such as an overly-dense crowd, and unusual crowd movement, occurs in the squares, **so that** I can further focus on the relative views, analyse them and determine a course of actions.	**As a Local Police officer:** - once being informed of anomalous events in crowded areas, I can activate appropriate actions such as dispatching the relevant authorities. - once there is an anomaly flagged, I am alerted and notified which live feed from which camera (the one capturing the location of the event) should be checked with attention and I should be able to view the feed a few minutes before the anomaly happened, to accurately assess the cause. - once I review the camera feed of the flagged anomaly, I have the option to verify if this event’s analysis is indeed correct or not. This option similar to GRN3 use case allows for continuously improving the system. - I also want to: i. view the anomalous events on the map, ii. compare events in different contexts (e.g., weekly market vs Christmas market), iii. view the timeline of events, iv. have a clustering of events, - I can prevent or be better prepared for a similar situation in the future.

4.1.4 MT3 – Monitoring of parking places

The target of this use case is the so-called “Ex Zuffo” Parking Area, which is one of the largest parking areas in Trento with around 1000 slots. It is used by the citizens and works well as an interchange car park, e.g., to leave the car and reach the city centre by public transportation, rentable bikes, and e-scooters.

To prevent robberies or damages to the cars parked, MARVEL Framework will support prevention activities with the audio-visual analysis of the existing cameras and the microphones that can be installed thanks to the MARVEL project. Furthermore, anomalous behaviours will be examined such as the correct use of parking spaces reserved for taxis, the occupation of spaces reserved for the vehicles of disabled people, the number of parked campers, and their time of stay, the average parking time of vehicles, the detection of possible damage, and other occurrences that will emerge during the execution of the experimentation. The audio-visual analysis must be carried out in real-time and on recording data saved on the servers of the Local Police.

The users of this system are intended to be Local Police officers. The User Stories for this use case are shown in Table 5 and give a better understanding of the functionality expected.

Table 5: MT3 – Monitoring of parking places

Title	Main user stories	Sub user stories and user intent
Monitoring of parking places, in terms of occupancy, proper use of the parking lots, and anomalous events.	**As a** Local Police officer, **I want to** monitor the use of the parking lot and detect an anomaly activity in parking spaces, **so that** I can take corrective /supporting actions.	**As a Local Police officer**: - in order to properly monitor the parking area, I want to view: i. the timeline distribution and duration of parking activity ii. the total number of vehicles iii. the clustering of vehicles and/or events iv. the severity and type of anomalies observed. - once an anomaly is flagged, I am alerted and notified which live feed from which camera should be checked with attention and I should be able to access the feed a few minutes before the anomaly happened, to accurately assess the event. - once I review the feed of the flagged anomaly, I have the option to mark this event as anomalous or not. This option similar to use case MT1 allows for continuously improving the system.

4.1.5 UNS1 – The drone experiment

The purpose of the drone experiment is to evaluate the potential of drones for monitoring large open-space public events. The UNS drone experiment setup includes computations resources, video recording from the drone and audio recording from the ground, which serves as a supporting modality for the inference. Such a setup could be useful for a quick security check in the case of crowded spaces. Commonly, street cameras can record only a frontal view of the event and inner details cannot be observed. Furthermore, there are angles or even whole spaces that are not covered using cameras. Flying over the crowd zones of interest, it can be quickly checked if there is some unusual problematic behaviour.

UNS is not a true pilot and there is not an end-user. User journeys, goals and requirements are set according to the discussion with the organisers of the EXIT festival – one of the largest

European music festivals that is being held in Novi Sad each summer. For R1, we will focus on crowd counting and prepare the AVCC component, whereas the full-scale version will include crowd behaviour classification. The user stories for this use case are shown in Table 6.

Table 6: UNS1 – The drone experiment

Title	Main user stories	Sub user stories and user intent
Crowd counting and anomaly detection	**As an** organiser of a large public event, **I want to**: - have an estimation of the number of people in the event, or a specific part of it (e.g., a stage) - be alerted if there is an overcrowded place and have real-time access to the corresponding AV stream for inspection and determining the corresponding course of actions.	**As an organiser of a large public event**, being aware of the number of visitors as well as their location at the area of the event but also being informed in real-time if there is an anomalous behavior in a part of the crowds is crucial for alerting people and relevant services. For such tasks, overcrowded places are of special interest due to difficult and slower access to emergency services. As it is difficult to detect anomalous events with ground-based cameras due to the limited visibility, I would like to have an option to have a live stream from drone-based videos, accompanied by audio-visual recordings from the ground. In such a way, decision-making could be accelerated.

5 Demonstration

In this section, we present the demonstrator that focuses on services provided from the end-user point of view. We describe the role of the DMT as a means for visualising the collected and processed relevant data and detected events. For each use case defined, we present the screens with which the end-user interacts.

5.1 Decision-Making Toolkit

5.1.1 GRN3 – Traffic Conditions and Anomalous Events

As described in Section 4.1.1, in this use case, traffic conditions are monitored to detect anomalous events, for example, traffic jams, accidents, cars stuck and obstructing a junction, very slow vehicles and service vehicles parked on the side or obstructing a carriageway. In Table 2, the user stories that are met through the DMT for this use case were described.

Users will firstly choose the use case of their interest from the available use cases, as shown in Figure 12.

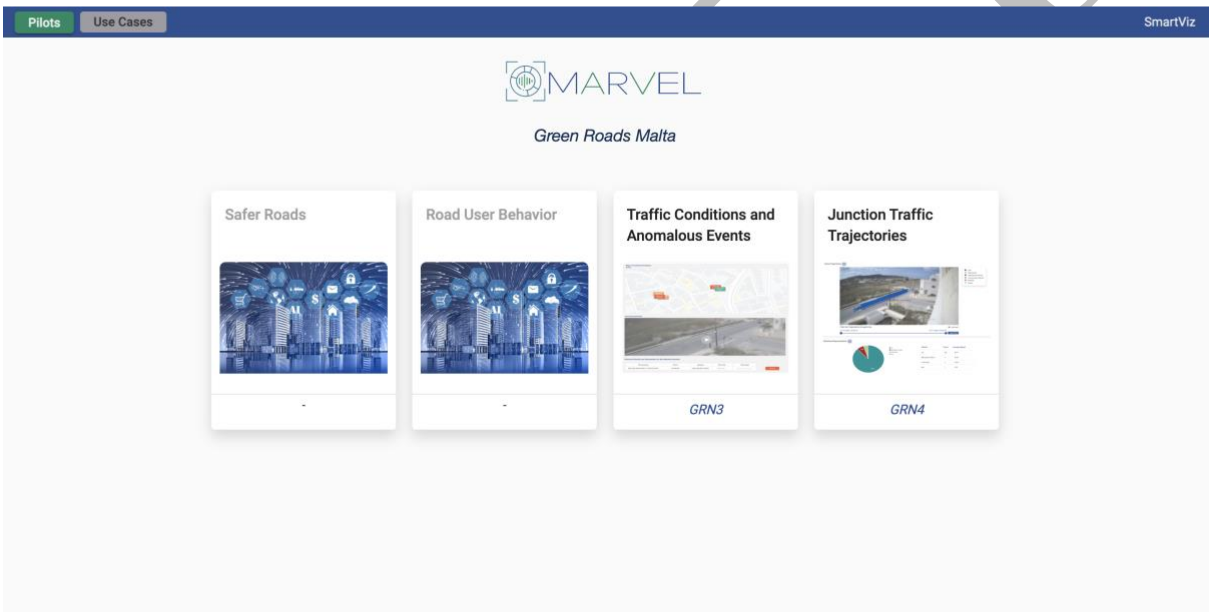

Figure 12: GRN Pilot Page

The dashboard for the GRN3 use case links to two implemented user stories: Anomaly detection and Traffic conditions. Both user stories are offered to a Traffic Manager as a user in a single dashboard depicted in Figure 13 to combine information and explore the available data and analysis. The goal of the anomaly detection user story is to alert the users if an anomaly occurs so they can analyse it and determine a course of actions.

Figure 13: GRN3 Dashboard

These detected anomalies are visualised in the Real-time Map Representation widget (see 3.3.7) in order to facilitate a quick response from the users' perspective so they can take appropriate actions (i.e., dispatching of the relevant civil protection authorities). The anomalies are detected by AVAD and served to SmartViz through Kafka messages by the Data Fusion Bus (DFB) component. The Video Player widget can either be used to play a live feed from the selected camera, or a static video that corresponds to a detected anomaly, upon user request.

The inference results and textual information linked to the detected anomalies are represented by the Details Widget (see 3.3.3). The data feeding this widget are historical and accessed through an Elastic search database which is constantly updated by the audio, visual, and multimodal AI subsystem components. Furthermore, SmartViz allows the users to validate the inference results by marking them, resulting in a continuous improvement of the systems and the AI models' training. The inference result verification is sent through Kafka messages by SmartViz to DFB which then updates the status of the corresponding event and stores them in an Elastic search index accessed by the Data Corpus.

SmartViz is a toolkit that also supports AV data representation. The Video Player widget (see 3.3.6) supports both real-time video streams from the cameras at the edge through an RTSP server source, and a static binary video output from StreamHandler that retrieves video snippets from the video streams that fall under the requested, by the user, time frame. Regardless the data source or the type of the AV streams, the user can take action and investigate the scene given multisource data combined in an intuitive way.

5.1.2 GRN4 – Junction Traffic Trajectory Collection

As described in Section 4.1.2, this use case is focused on the requirement of long-term data analytics that shed light on both the behaviour of road users (e.g., car drivers, motorcyclists, cyclists, pedestrians, etc.) and on gathering traffic statistics on road network junctions. In Table 3, the user stories that are met through the DMT for this use case were described. Hence, Figure 14 is the dashboard that combines the two user stories and contains all the widgets and visualisation schemas as presented below.

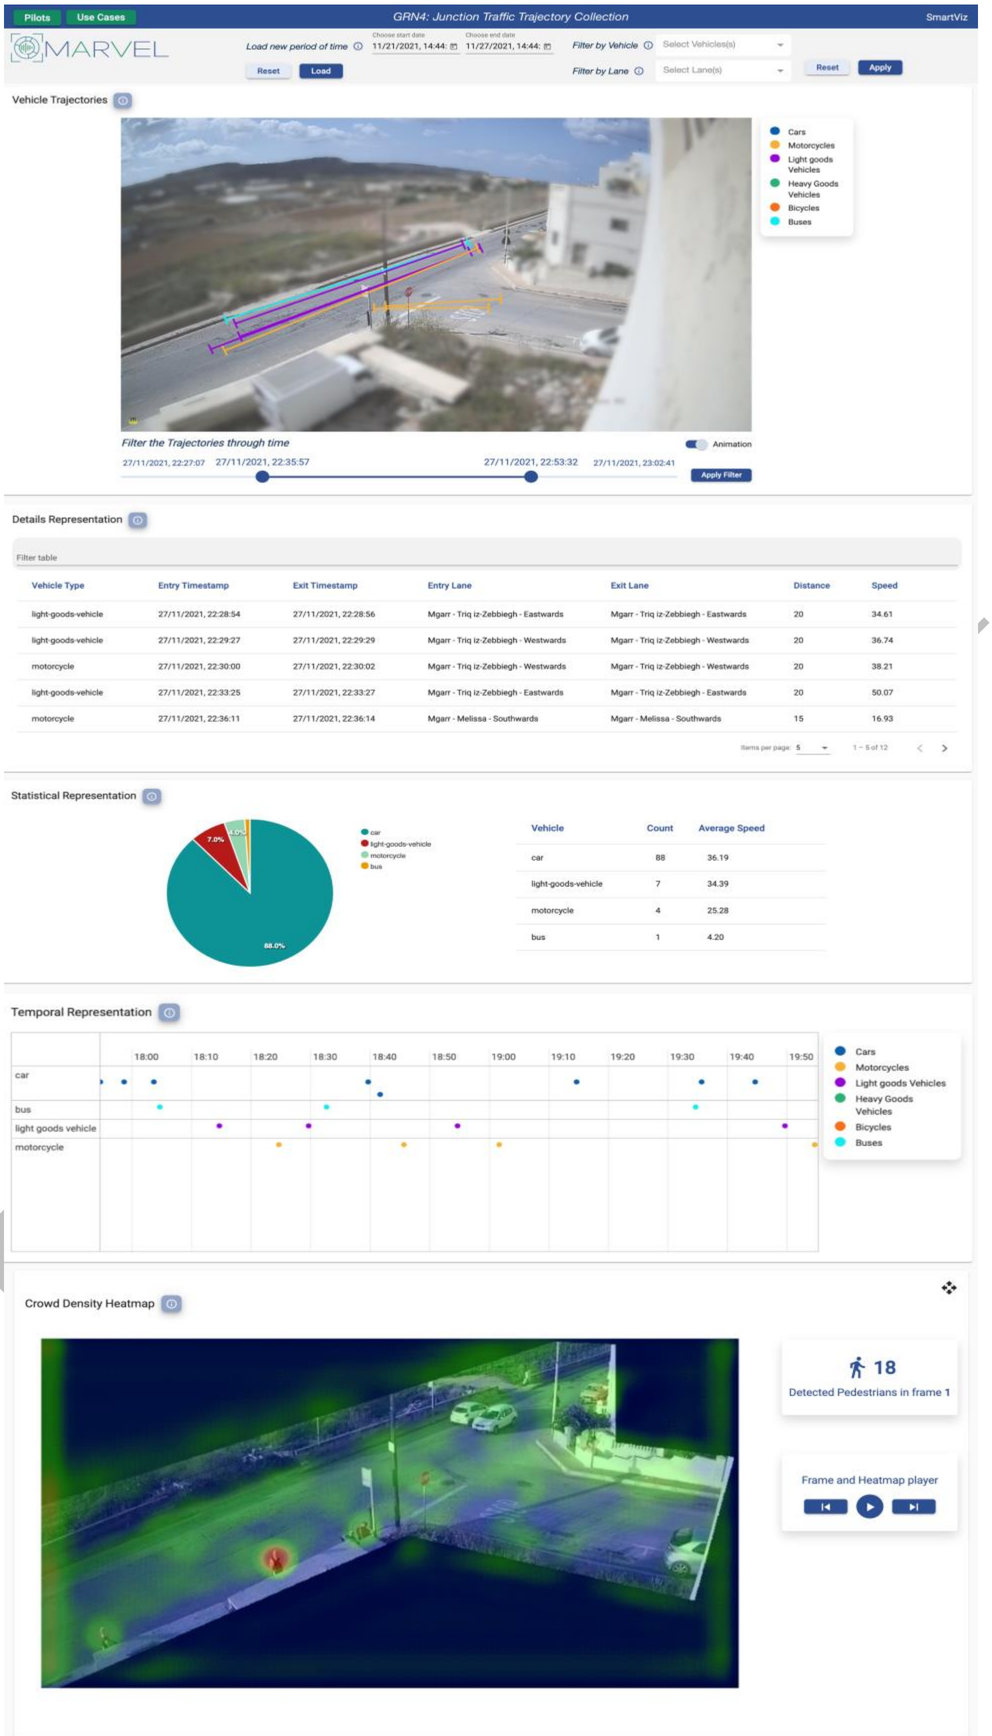

Figure 14: GRN4 Dashboard

This use case will allow traffic engineers to make informed decisions on issues such as the most preferred location of a cycle path or if there needs to be more enforcement on issues such as heavyweight vehicles staying on the left side of the road. This use case was also selected for the MVP release at M12.

To illustrate the trajectories of the detected vehicles and detailed information, we developed a visualisation widget to be added to our pool of widgets, the Vehicle Trajectories widget (see 3.3.4). In this widget, the paths of the passing vehicles that are detected by the CATflow component are drawn in the image of the camera feed that is recording them. The paths are grouped and colour-coded depicting different types of vehicles. The users are also able to change the time period and filter by the type of vehicle or lane to further investigate.

Furthermore, detailed available information about the detected incoming events is represented in a tabular form in the Details Widget (see 3.3.3). This widget also supports a standalone text filtering capability in order to facilitate a rapid search through the available information.

Additionally, this use case presents a count (a distribution) of the different types of vehicles passing through a road junction. This will allow the user to be able to track how frequently a road is being used and by what type of vehicle which could then be used to make data-driven decisions on the road infrastructure. For example, a road frequently used by heavy vehicles would need to be re-surfaced more often and/or would require a different mix of road surfacing materials. Alternatively, the engineer might consider rerouting the heavy vehicles. On the other hand, a road frequently used by cyclists would make a candidate for the installation of a cycling lane or the installation of cycling infrastructural technology. As a result, simple data processing algorithms are implemented such as collecting and separating the traffic counts based on the different types of vehicles or the visualisation of colour-coded trajectories which a human being would be able to visually process and reach conclusions.

Moreover, the information of detected vehicles that is the outcome of the Sound Event Detection (SED) component is also presented to the user in a temporal form, through the Temporal Representation widget (see 3.3.1).

Finally, the AVCC component outputs a JSON with points on each frame of a video with a likelihood value for a pedestrian to be present at that point. This information is summarised in a matrix that corresponds to the input frame, and it is later visualised as a heatmap over the original frames in SmartViz through the Crowd Density Heatmap widget (see 3.3.2).

Since the MVP, there were additions to the filtering options and the interactions of the user with the dashboard (i.e., download datasets), the combination of the two different dashboards for each scenario into one complete dashboard similar to all the other selected use cases. Frames, images, and annotations were added on top of them, from different hours of the day to the Vehicle Trajectories and Crowd Density Heatmap widgets, to better enhance the user experience and data analysis and exploration.

5.1.3 MT1 – Monitoring of Crowded Areas

As described in Section 4.1.3, this use case is focused on the selection of views of areas of interest for monitoring situations such as exceptional crowd, suspect or unusual crowd movements, etc. In Table 4, the user story that is met through the DMT for this use case was described. First, users will be able to use the desired use case from the available use cases, as shown in Figure 15.

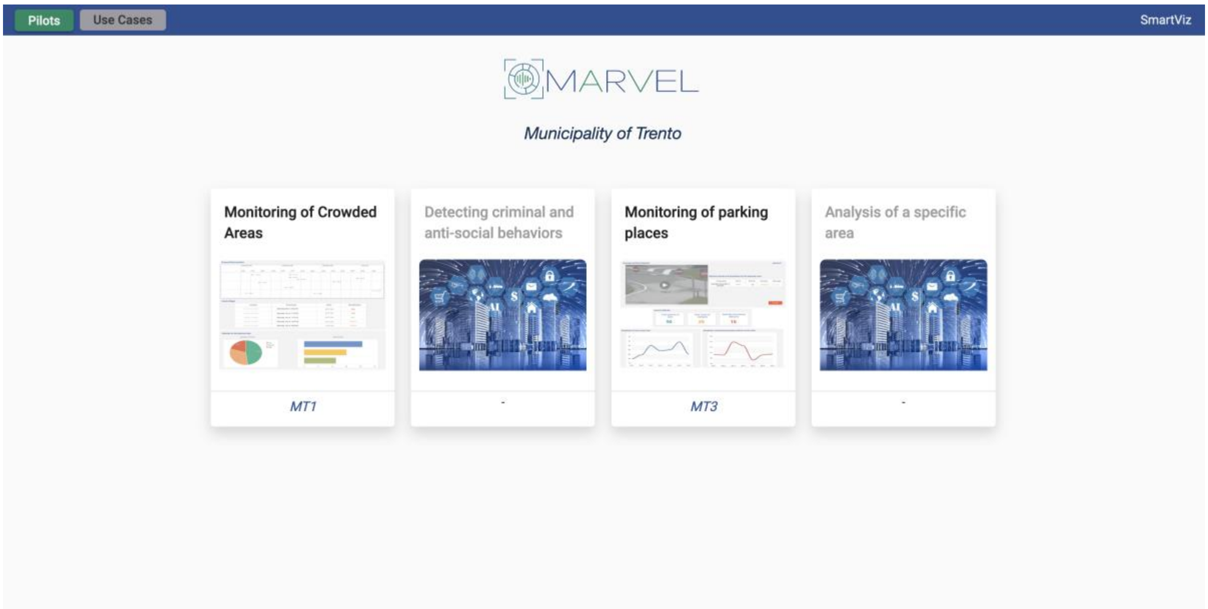

Figure 15: MT Pilot Page

In Figure 16, the constructed dashboard containing the widgets, filters, and features described below for the Monitoring of Crowded Areas use case is represented.

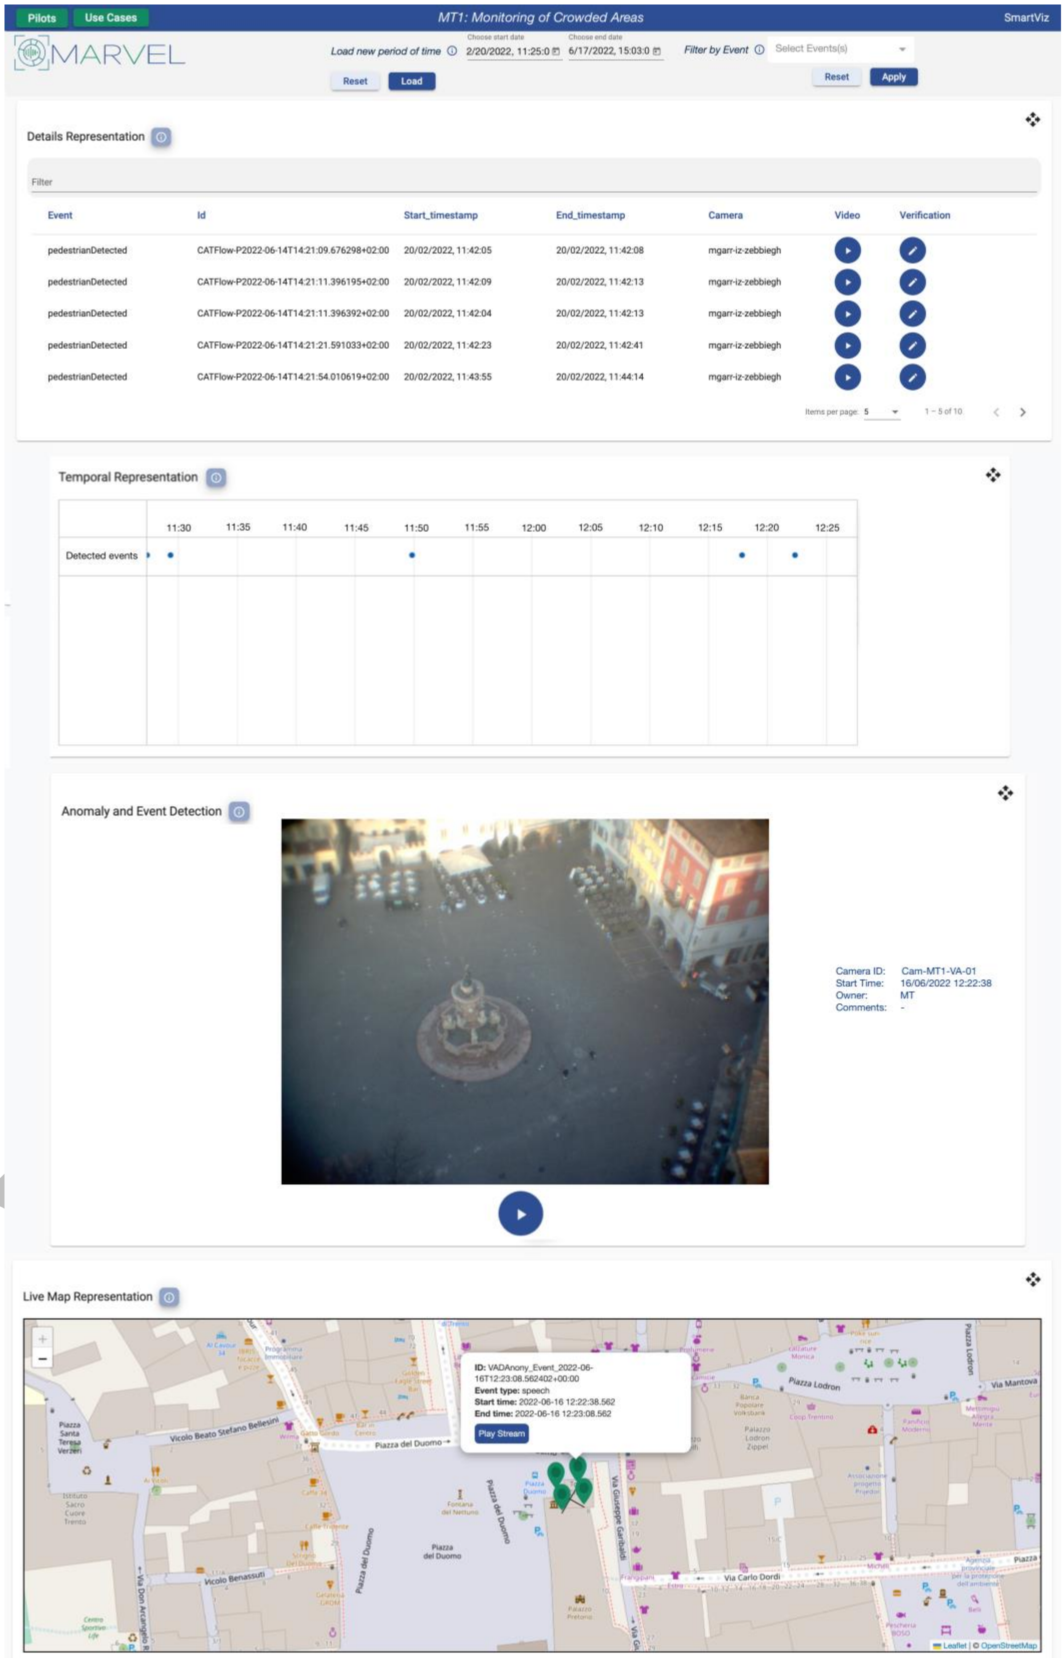

Figure 16: MT1 Dashboard

The users of this use case are local police officers that want to be alerted if an anomaly such as an overly dense crowd, and unusual crowd movement, occurs in the squares, so they can further focus on the relative views, analyse them and determine a sequence of actions. In order to notify the users of the detected alerts, their location and any available information, we chose the Real-time Map Representation widget (see 3.3.7). Hence, the detected alerts are represented as pins in a map in real-time. The users are also given the opportunity to see either the video stream of the camera, a video corresponding to an event, or a video that occurred in a selected time frame through the Video Player widget (see 3.3.6).

The inference results and all the available information linked to the detected alerts are also represented in the Details Widget (see 3.3.3), where the user can also verify their validity also in the Temporal Representation widget (see 3.3.1).

5.1.4 MT3 – Monitoring of Parking Places

As was described in Section 4.1.4, this use case is focused on the prevention of robberies or damages to vehicles parked in a monitored parking place, the examination of anomalous behaviours and events such as the incorrect use of parking spaces, the occupation of spaces reserved for the vehicles of disabled people etc. In Table 5, the user stories that are met through the DMT for this use case were presented. In Figure 17, the constructed dashboard containing the widgets, the filters and the features described below for the Monitoring of Parking Places use case is represented.

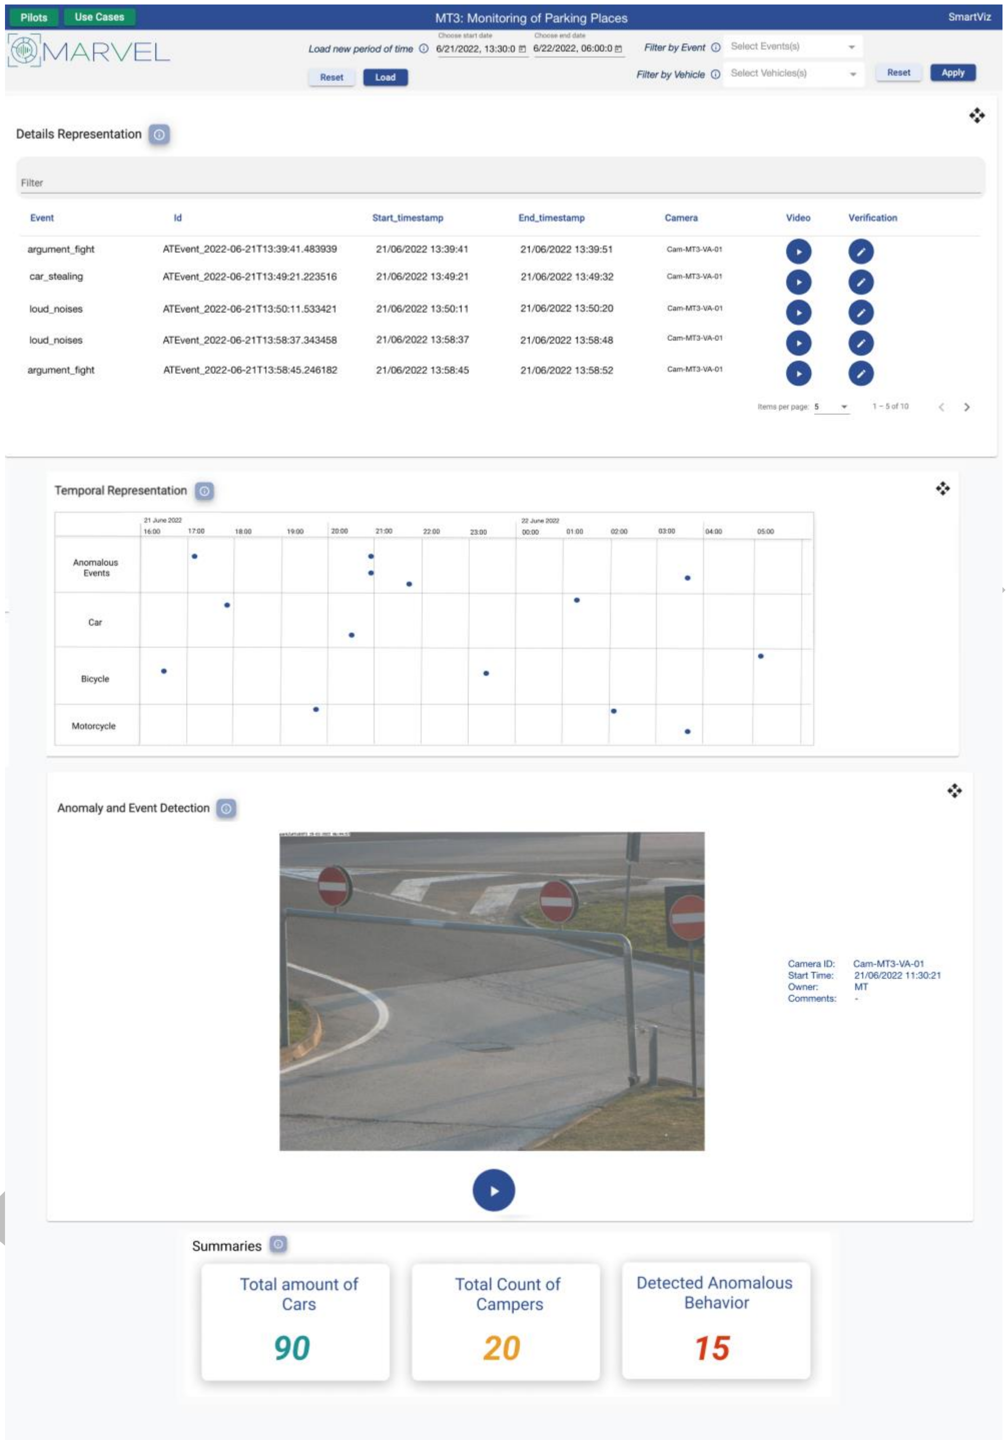

Figure 17: MT3 Dashboard

In this particular use case, four visualisation widgets carefully selected by SmartVizs' pool of widgets, are being utilised. First of all, the inference results and all the available information linked to the detected vehicles and anomalous events by AT and SED are represented in the Details Widget (see 3.3.3) where the user can also verify their validity, and also in the Temporal Representation widget (see 3.3.1). Hence, a temporal analysis of the historical data can be performed by the user. The user has also the opportunity to see either the video stream of the camera, a video corresponding to an event, or a video that occurred in a selected time frame through the Video Player widget (see 3.3.6). Finally, information regarding the total number of each vehicle type detected in a parking lot and the detected anomalies or behaviours are represented in the Summaries widget (see 3.3.8).

5.1.5 UNS1 – Drone experiment

As was described in Section 4.1.5, the purpose of this use case is to evaluate the potential of drones for monitoring large open-space public events. The UNS drone experiment setup includes computations resources, video recording from the drone and audio recording from the ground, which serves as a supporting modality for the inference. Such a setup could be useful for a quick security check in the case of crowded spaces. In Table 6, the user stories that are met through the DMT for this use case were described. As a first step, the users are able to choose the use case of their interest from the available use cases of the UNS pilot dashboard, as shown in Figure 18.

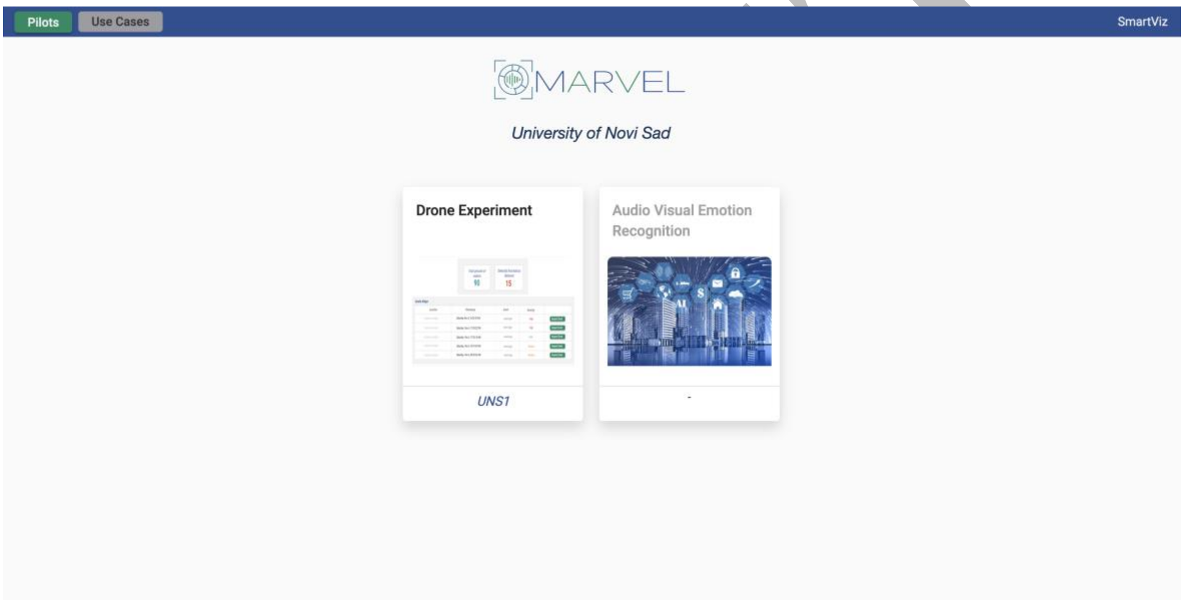

Figure 18: UNS Pilot page

In Figure 19, the constructed dashboard containing the widgets, the filters, and the features for the Drone Experiment use case is illustrated.

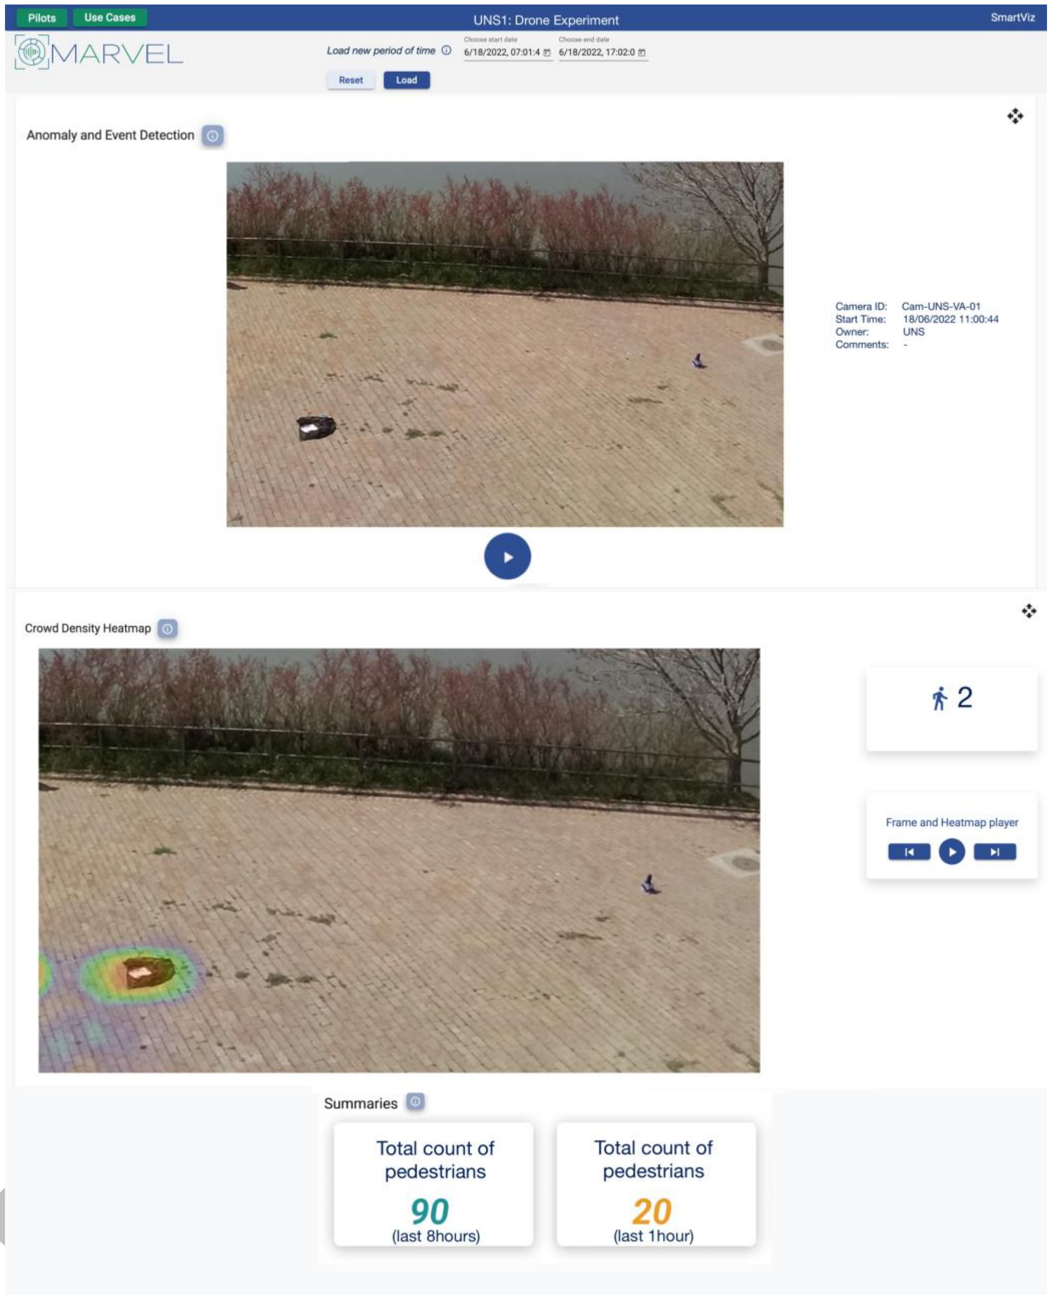

Figure 19: UNS1 Dashboard

The end-user of this use case is an organiser of an event that wants to have an estimation of the number of people in the event or a specific part of it, and furthermore, to be alerted if there is an overcrowded place, to have real-time access to the corresponding AV stream for inspection, and to determine the sequence of actions they should follow. The first requirement is met by the Summaries widget (see 3.3.8) which shows the total number of visitors and the detected anomalous events in a selected time frame. This requirement is also met by the Crowd Density Heatmap widget (see 3.3.2) that presents the probability of detected pedestrians being present in a video frame. Both visualisation widgets are having as input the output analysis of the AVCC component. The latter requirement from the user is met by the Video Player widget (see 3.3.6) where it displays the static video of a detected event or live camera stream.

6 Contribution to MARVEL goals

6.1 Project-related KPIs

The work performed in this task directly addresses the project-related KPI-O3-E4-1 *"Detailed insights into more than 5 hidden correlations"*. This KPI is naturally linked to Enabler 4 *"Complex decision-making and insights"* of objective 3 of the project: *"Break technological silos, converge very diverse and novel engineering paradigms and establish a distributed and secure Edge-to-Fog-to-Cloud (E2F2C) ubiquitous computing framework in the big data value chain."* However, T4.2, T3.3 as well as the whole WP6 where the real-life experiments are being performed are of key importance, since they feed the visualisations (T3.3 and T4.2) and evaluate them (WP6) respectively. This KPI was more closely monitored after the MVP release, where end-users had the opportunity to experiment with the MVP version of the DMT, based on the hypotheses already being investigated, regarding where to find hidden correlations. After the MVP release and the end-users' feedback as well as a few brainstorming discussions, we have now decided, that the hypotheses to explore further will be the following:

(i) *Investigate the role of time and weather in the appearance of traffic anomalies as well as in the appearance of unusual events in the use cases related to public safety*. A preliminary implementation of this investigation can already be seen in this initial version of the DMT where the weather widget is added to the dashboard.

(ii) *Investigate what happens at the same period of time across different parts of the city based on historical data.* This hypothesis is at the beginning of its implementation with the option of split screens, fed by different "data batches" being one of the techniques we have been experimenting with.

(iii) *Correlate the number of vehicles at a road or a junction and investigate the possibility of alerts for maintenance or repair needs*. This hypothesis is still under investigation. After the release of the first integrated version of the MARVEL framework, we are planning to interview the end-users of the first version of the DMT and explore types of thresholds for alerts related to the topic and ways for them to be visualised.

Each of these hypotheses can lead to more than 1 hidden correlation or even to more hypotheses after the first integrated version of the MARVEL framework will be released and the users will have the chance to play with it. Progress will be monitored regularly.

6.2 Component-related KPIs

As already mentioned, SmartViz is the key component used for the development of the DMT. The baseline measures of the component-related KPIs, namely usability, scalability, availability, reliability, and performance were initially measured after the release of the MVP at M12. The results can be seen in detail in Deliverable D5.2[11]. In short, the MVP scored very well in terms of usability (top10%). Unfortunately, availability presented only 91,7% instead of the 100% target. The root cause investigation spotted the issue on the hosting component which should not cause any further incidents during and after the release of M18. The reliability of what the tool was showing and the input it was receiving was based on internal checks and was 100% as per target. Given the small scale of the MVP scalability and performance were not properly tested, however, ZELUS plans to be supporting Big Data without errors that hinder

[11]"D5.2: Technical evaluation and progress against benchmarks – initial version, 2021. https://doi.org/10.5281/zenodo.6322699

the usability and performance of the interface, allowing a maximum of 8-10 seconds response time.

Moving a few months forward to today and the upcoming release of the first integrated version of the MARVEL framework, there is already a number of improvements related to SmartViz since its MVP edition. These improvements are going to be validated on a wide scale after M18 when a larger number of end-users will have the opportunity to test it and give their feedback.

It is important to mention here that other than the already known components KPIs, qualitative component upgrades include the extension of the features, widgets, filtering options, and the overall capabilities of SmartViz since the beginning of the project as well as since the MVP, More specifically, ZELUS has developed four new widgets since the beginning of the project until M18 solely for the user requirements and visualisation needs of the MARVEL AI models outputs. The Real-time Crowd Density Heatmap, the Vehicle Trajectories, the Map Representation, and the Video Stream Player widgets were developed to support and represent data for detected anomalies, events, and the depiction of moving entities i.e., vehicles, pedestrians, in monitored areas i.e., junctions, city squares, and their behavioral process. Equally important for enriching SmartViz's pool of widgets are the filtering options we adjusted and set for the meaningful analysis, investigation, and data exploration.

7 Conclusions

This deliverable presents the methodology for the design of the DMT and its building elements. More specifically, the document presented the outcomes of the involvement of end-users in the initial stages of the development activities of MARVEL DMT. Having described the UCD process, we proceeded with an overview of the toolkit functionalities and the definition of the main use cases, the user journeys, and the intent behind each functionality. The tool comprising the DMT (SmartViz) has been presented in detail and its use inside the platform has been set.

Based on a well-defined UCD process that aims to achieve a high degree of usability for the MARVEL DMT, the initial version of the toolkit has been demonstrated and includes the functionalities required to support the selected use cases.

The next steps of this work include the continuous evaluation of the MARVEL framework including the DMT and the validation of its functionalities in real-world scenarios. The evaluation process already started right after the delivery of the MVP in M12 and it will be continued after the release of the initial version of the integrated MARVEL framework. This validation will allow us to further align with the users' expectations and adjust the future development activities. Furthermore, focused evaluation on the usability will reveal any possible early deviations from what the users expect from the DMT, thus allowing us to perform corrective actions and to update the way that DMT's capabilities are presented to the end-users in view of the final release.

8 References

Jeffrey Rubin, D. C. (May 2008). *Handbook of Usability Testing: How to Plan, Design, and Conduct Effective Tests* (2nd ed.). Wiley ed.