AI4Gov

Trusted AI for Transparent Public Governance
fostering Democratic Values

Deliverable 4.4

Policies Visualization Services V2

<31-03-2025>

Version 1.0

Funded by the European Union	*Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the Agency. Neither the European Union nor the granting authority can be held responsible for them.*

PROPERTIES	
Dissemination level	Public
Version	1.0
Status	Final
Beneficiary	UPRC
License	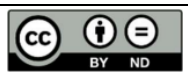 This work is licensed under a Creative Commons Attribution-No Derivatives 4.0 International License (CC BY-ND 4.0). See: https://creativecommons.org/licenses/by-nd/4.0/

AUTHORS		
	Name	**Organisation**
Document leader	Andreas Karabetian	UPRC
Participants	Dimitris Apostolopoulos	VVV
	Alenka Guček	JSI
	Xanthi Papageorgiou	UBI
	Celia Parralejo	DPB
Reviewers	Theodore Chadjipadelis	AUTH
	Georgia Panagiotidou	AUTH
	Vicky Stroumpou	MT
	Stavroula Kefala	MT

VERSION HISTORY				
Version	**Date**	**Author**	**Organisation**	**Description**
0.0	15/02/25	Andreas Karabetian	UPRC	Initial ToC
0.1	17/02/25	Andreas Karabetian	UPRC	Updated Abstract and Conclusions
0.2	27/02/25	Andreas Karabetian	UPRC	Updated Chapters 1 and 2
0.3	11/03/25	Dimitris Apostolopoulos	VVV	Updated Chapter 7
0.4	21/03/25	Xanthi Papageorgiou	UBI	Updated Chapter 3
0.5	22/03/25	Alenka Guček	JSI	Updated Chapter 4
0.6	24/03/25	Alenka Guček	JSI	Updated Chapter 6
0.7	24/03/25	Celia Parralejo	DPB	Updated Chapter 5
0.8	24/03/25	Andreas Karabetian	UPRC	Preparation for internal review
1.0	31/03/25	Spiros Borotis	MAG	Final review and submission

Table of Contents

List of figures

List of Tables

Abbreviations

Abbreviation	Description
AI	Artificial Intelligence
API	Application Programming Interface
OECD	Organisation for Economic Cooperation and Development
SDG	Sustainable Development Goals
VVV	Vari-Voula-Vouliagmeni
DPB	Diputación Provincial de Badajoz
IRCAI	International Research Centre On Artificial Intelligence
UNESCO	United Nations Educational, Scientific and Cultural Organization
KPI	Key Performance Indicator
ML	Machine Learning
PRO	Patient-reported outcomes
CRO	Caregiver-reported outcomes

Abstract

The Visualization Workbench, the evolved central frontend component of the AI4Gov project, builds upon its transformative foundation to deliver a more powerful and comprehensive platform for data-driven governance. This new iteration expands its capabilities with enhanced functionalities, improved multilingual support, and a broader spectrum of visualization tools. A key advancement in this advanced version of the Workbench is its newly introduced mobile-friendly design, ensuring seamless access across devices—a feature that was previously unavailable. Now encompassing a richer suite of analytical and interactive features, the Workbench empowers policymakers with deeper insights into complex documents, fostering more effective and transparent decision-making. The integration of advanced AI-driven visualization techniques further refines its role in Sustainable Development, Tourism, and Water Management, offering an even more nuanced understanding of sector-specific challenges. The Workbench enhances its core services, including the Bias Detector Toolkit and the Policy Recommendation Toolkit, with more sophisticated visualization layers that make biases and policy impacts even more interpretable. The improvements in usability and accessibility ensure that stakeholders—ranging from policymakers to researchers—can seamlessly engage with the data, uncover hidden patterns, and validate AI-driven insights with greater clarity. By strengthening its commitment to ethical AI deployment, transparency, and accountability, the Workbench continues to serve as a pivotal tool in shaping responsible AI governance, ensuring fairness, equity, and societal well-being.

1 Introduction

1.1 Purpose and scope

The purpose of this Deliverable is to present the next evolution of Visualization Workbench, this dynamic web application, which serves as the centralized platform for an array of innovative solutions at the intersection of artificial intelligence and governance. As the project's front end, it continues to exemplify the convergence of cutting-edge technology and real-world challenges, now with expanded functionalities, improved multilingual support, and a fully mobile-friendly design. Building upon its strong foundation, it enhances its role as the hosting ground for all pilot partner use cases, seamlessly integrating the collective expertise of three visionary partners. These partners, strategically active in Water Management, Sustainable Development, and Tourism, contribute domain-specific knowledge that enriches the platform's ability to address sector-specific challenges. With a broader array of visualization tools and deeper analytical capabilities, the Workbench empowers stakeholders with even more intuitive and actionable insights.

The Workbench now accommodates an expanded and more refined set of use cases, further unlocking the synergies between AI and governance. Beyond its role as a showcase for pilot projects, it extends its impact through enhanced versions of its core services, the Bias Detector Toolkit and the Policy Recommendation Toolkit, which now offer more sophisticated visualization layers. These improvements reinforce fairness, transparency, and informed decision-making, ensuring that AI-driven governance continues to align with ethical and responsible practices.

1.2 Document structure

The deliverable is structured as follows:

- **Chapter 1** introduces the deliverable's purpose and scope, it's document structure along with any updates compared to the previous version.
- **Chapter 2** presents the updated version of the Visualization Workbench component, the direct result of Task 4.4, focusing on a state-of-the-art analysis, an architecture diagram, explaining the baseline technologies, source code and availability as well as installation instructions.
- **Chapter 3** describes the updated Policy Recommendation Dashboard, the front-end application of the Policy Recommendation Toolkit. In this chapter we take a brief overview of the component's updated interface, it's key objectives as well as changes to the user interaction.
- **Chapter 4** provides information on the updated Bias Detector Toolkit Dashboard, with a brief overview of the component's updated interfaces, the key objectives as well as changes to the user interaction.

- **Chapter 5** focuses on the first sector of use cases, Water Management. Then it delves deeper into each updated use case accompanied by useful screenshots.
- **Chapter 6** focuses on the second sector of use cases, Sustainable Development. Then it delves deeper into each updated use case description accompanied by useful screenshots.
- **Chapter 7** focuses on the third sector of use cases, Tourism-driven Policy Management and Optimisation. Then it delves deeper into each updated use case description accompanied by useful screenshots.
- **Chapter 8** concludes this deliverable, summarising the key findings.
- **Chapter 9** includes the reference list.

1.3 Updates with respect to previous version

The latest iteration of the Visualization Workbench represents an evolution from its previous version, incorporating a host of improvements aimed at refining usability, accessibility, and analytical depth. One of the most prominent enhancements is the introduction of a fully mobile-friendly design, which was previously absent, thereby ensuring seamless cross-device functionality and expanding the Workbench's accessibility to a wider range of users. The component's spectrum has been significantly enriched with new AI-driven techniques, enabling policymakers to explore complex datasets with greater granularity and clarity, particularly in the domains of Sustainable Development, Tourism, and Water Management. Moreover, the Bias Detector Toolkit and the Policy Recommendation Toolkit have undergone notable upgrades, with more sophisticated visualization layers that enhance interpretability and facilitate a clearer understanding of biases and policy impacts. Multilingual support has also been significantly improved, reinforcing the Workbench's role in diverse governance environments by allowing non-English-speaking stakeholders to engage more effectively with AI-driven insights.

All these significant changes have been documented in the following chapters, thus enhancing the first version of the Deliverable by providing a more comprehensive and up-to-date analysis of the platform's advancements.

2 Visualization Workbench

2.1 Visualization tools for Policy Making and Bias Detection

The Visualization Workbench Component represents a pioneering effort at the intersection of artificial intelligence and governance, evolving into a more powerful and comprehensive platform that advances the role of data visualization in decision-making. The integration of advanced visualization tools into governance frameworks has become increasingly critical for enhancing policy formulation, transparency, and accountability. The sustainable development and deployment of AI in the public sector requires ongoing deliberation among developers, decision-makers, implementers, end users, and the public (leikas, 2022). The long-term success of AI-driven solutions in governance depends on their seamless integration into policy implementation frameworks (Mikhaylov, 2018), underscoring the necessity of collaboration between academia, government institutions, and the private sector.

With this in mind, the latest iteration of the Visualization Workbench builds upon its foundational capabilities to deliver a more robust and user-centric interface that empowers decision-makers across diverse sectors. Its enhanced multilingual support expanded suite of AI-driven visualizations, and mobile-friendly design mark a significant step forward in making AI-assisted governance more accessible and efficient. Additionally, the importance of bias detection tools in decision-support systems remains crucial, as highlighted by (Ltifi, 2016), who proposed visualization techniques for intelligent decision-support systems to analyze large-scale datasets. Similarly, (Law, 2020) emphasized the financial and ethical risks associated with undetected biases in machine learning, a challenge that the Visualization Workbench directly addresses through its upgraded Bias Detector Toolkit. This refined feature enhances the interpretability of AI-driven policy insights, reinforcing the commitment to ethical and fair AI governance.

Moreover, (Wu, 2020) demonstrated the application of interactive web-based geovisual analytics platforms for co-clustering spatiotemporal data, enabling simultaneous analysis of spatial patterns—a capability that the Workbench now supports more effectively through its improved analytical and interactive visualization features. As the state-of-the-art landscape continues to evolve, the Visualization Workbench stands not just as a tool but as an adaptive and dynamic response to the growing complexities of data-driven governance. With the advancements documented in this deliverable, it further solidifies its role as an essential component in modern governance, bridging cutting-edge research with practical policymaking needs.

2.2 Architecture

An image of the updated component architecture can be seen below, in Figure 1.

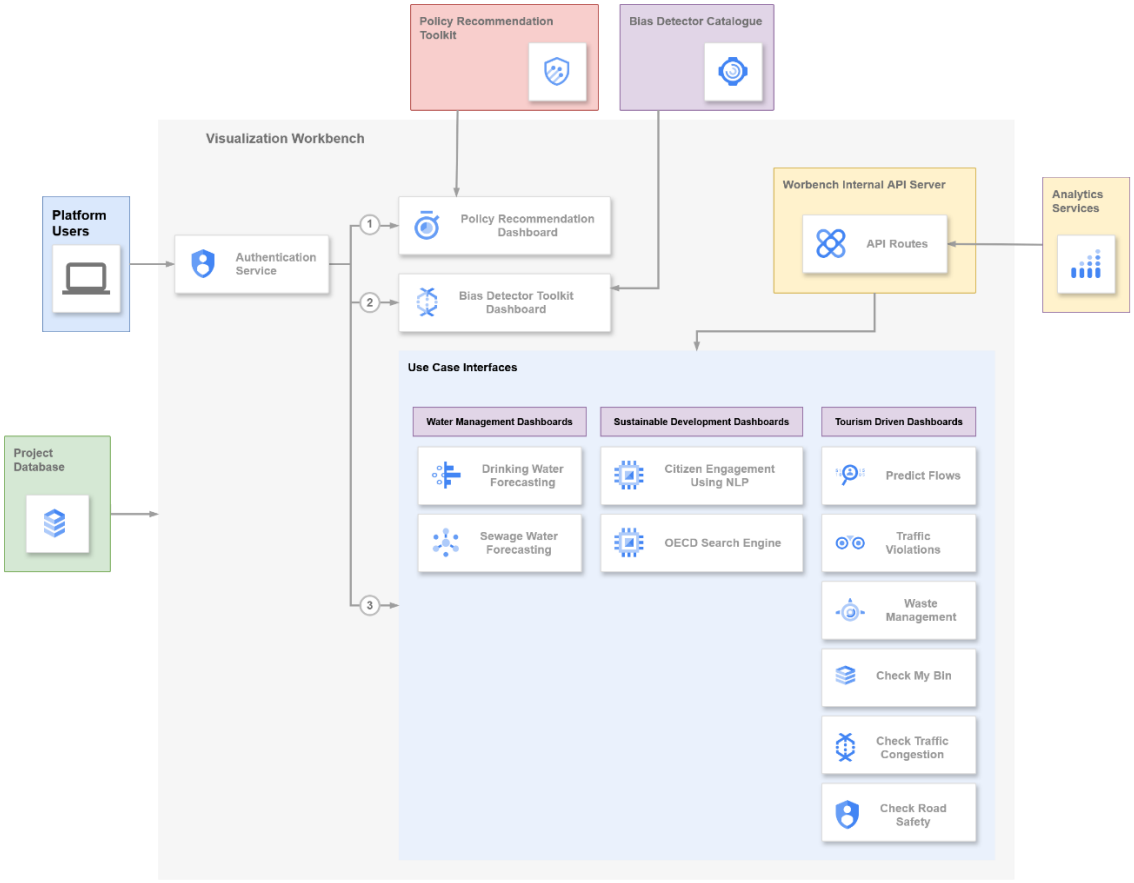

Figure 1: Visualization Workbench Architecture

2.3 Baseline Technologies

At its core, the Workbench is built as a Next.js web application, harnessing the power of this React based framework to create a dynamic and responsive user interface. This choice facilitates the rapid development of interactive and visually appealing features essential for effective data visualization.

Complementing the front end, the Workbench relies on an internal API server, leveraging the asynchronous, event-driven architecture to efficiently handle data requests from any other component within the project's architecture. The use of Docker (Docker, n.d.) adds a layer of versatility and consistency to the deployment process by containerizing the Next.js full-stack application.

This containerization ensures that the Workbench component, with all its dependencies, can be seamlessly packaged and executed in parallel with all the other components in the Kubernetes (Kubernetes, n.d.) cluster of the project. Together, these baseline technologies form a well-

integrated and agile foundation for the Visualization Workbench, offering a harmonious blend of front-end sophistication, server-side responsiveness, and deployment flexibility.

2.4 Source Code and Availability

The source code is available on the project's GitLab repository under the project named "T4.4 - Visualization Workbench", as seen in Figure 2.

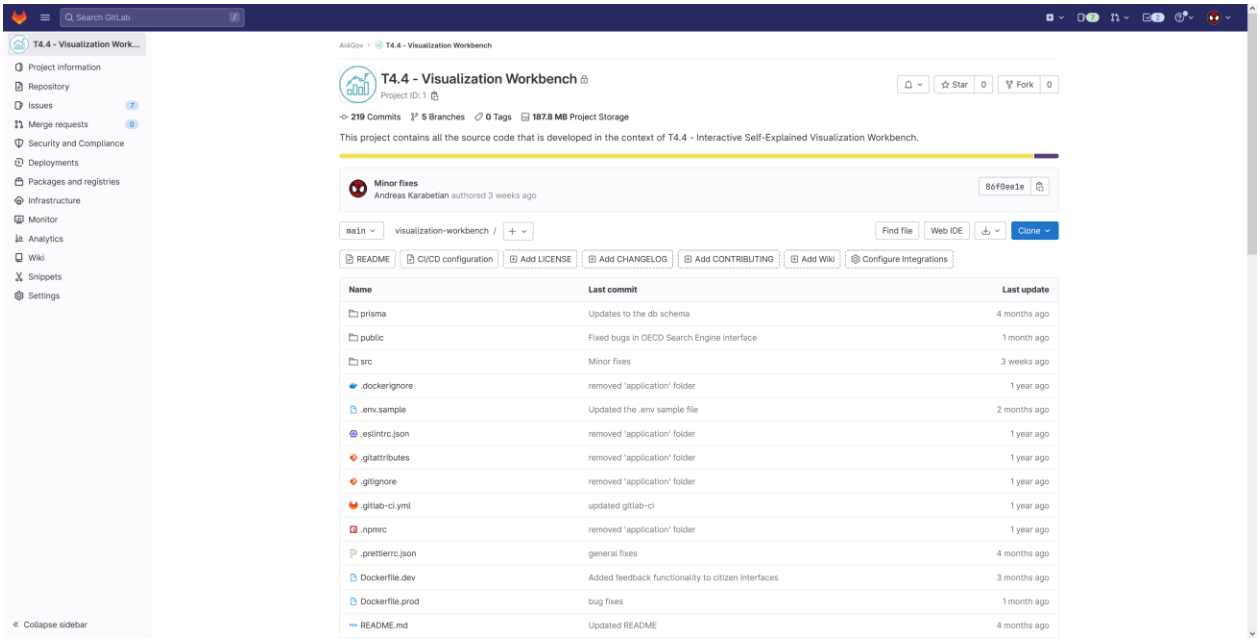

Figure 2: Visualization Workbench source code

In order to keep track of all use case requirements and maintain an organized repository, the Visualization Workbench utilizes GitLab issues and issue boards as seen in Figure 3.

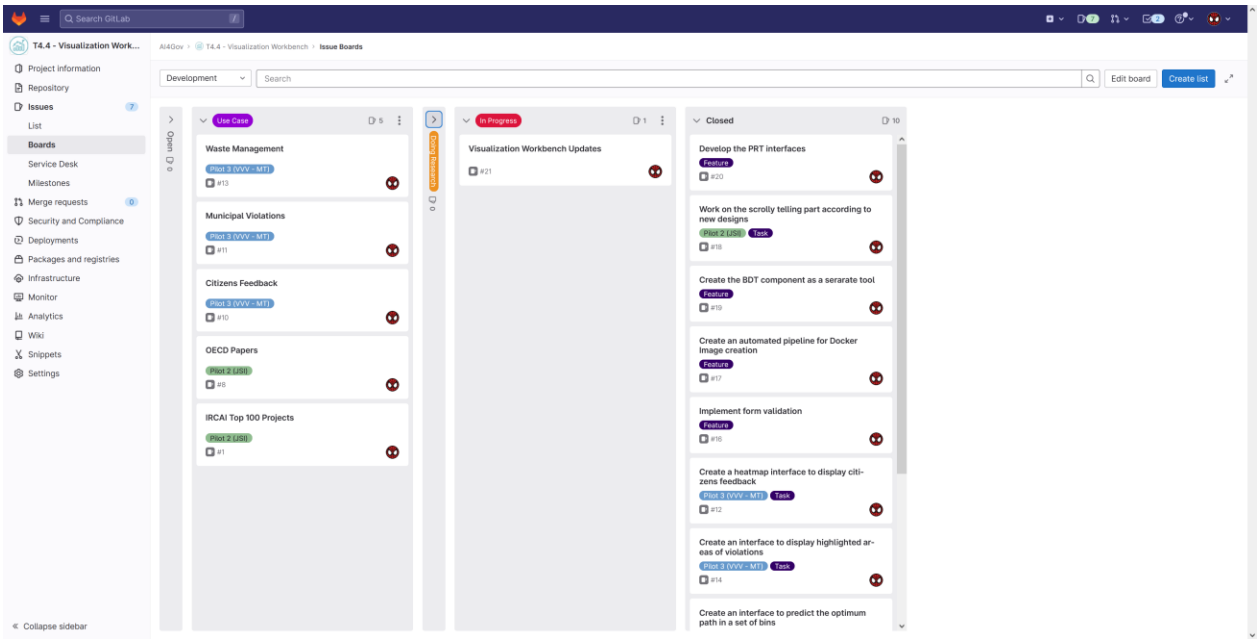

Figure 3: GitLab issues for Visualization Workbench project

2.5 Installation Guide

The installation of the component can be done using any container management solution, as it is fully containerized and the image is available through the AI4Gov GitLab container registry. Full instructions are also included in the repository README file. Once the container starts, the front-end can be found on port 3000.

3 Policy Recommendation Dashboard

3.1 Overview

The Policy Recommendation Toolkit aims at aiding end-users in finding, designing and comparing policies with reference to specific KPIs. It has a front-end component for user-interaction and a backend that also connects to the blockchain for policy administration and for defining and executing custom policies that can be created via common business rules established in smart contracts.

The Policy Recommendation toolkit can be leveraged by the Visualization Workbench component by integrating its functionality and linking it to menu action items that can then render its front-end interface within the workbench. Being deployed as a microservice and using the Next.js framework, the Policy Recommendation Toolkit can be integrated either by loose coupling or by directly integrating its source code into the relevant parts of the Visualization Workbench.

The Policy Recommendation toolkit is described in details in the D3.4 – Policy Recommendation Toolkit V2, which is due on M27.

3.2 Policy Recommendation Toolkit Interface

3.2.1 Overview

The interface of the framework includes interactive graphical components for providing each one of the major functionalities that fulfil the framework's objectives. There is more than one navigation option among the available tools through the application's overview screen, which are available to the policy maker; this can be chosen via the initial page (Figure 4 and Figure 5). As shown in the Figure 4, the four main functionalities are described, allowing the user to navigate to one of them. On the same screen, there are four categories in which the policy maker can create a new policy, as depicted in the Figure 5.

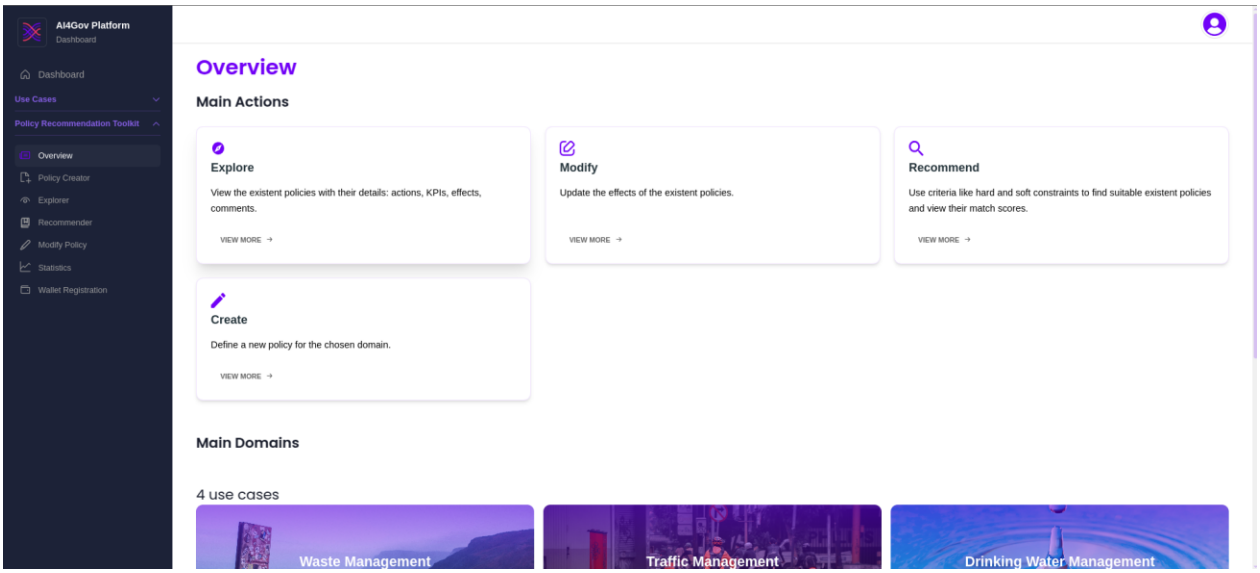

Figure 4: Homepage of the Policy Recommendation Toolkit (top view with functionalities)

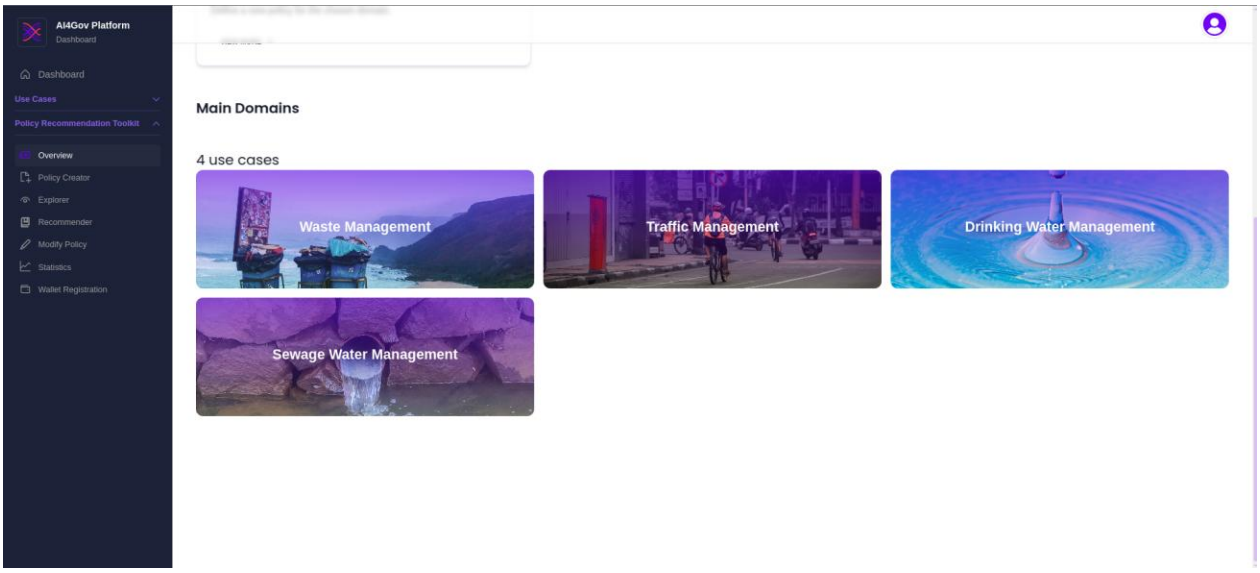

Figure 5: Homepage of the Policy Recommendation Toolkit (bottom view with the categories)

Then the user can create a policy from the relative menu, selecting a specific category (waste management, traffic management, drinking water management, sewage water management), Figure 6. The specific fields are described in detail to D3.4.

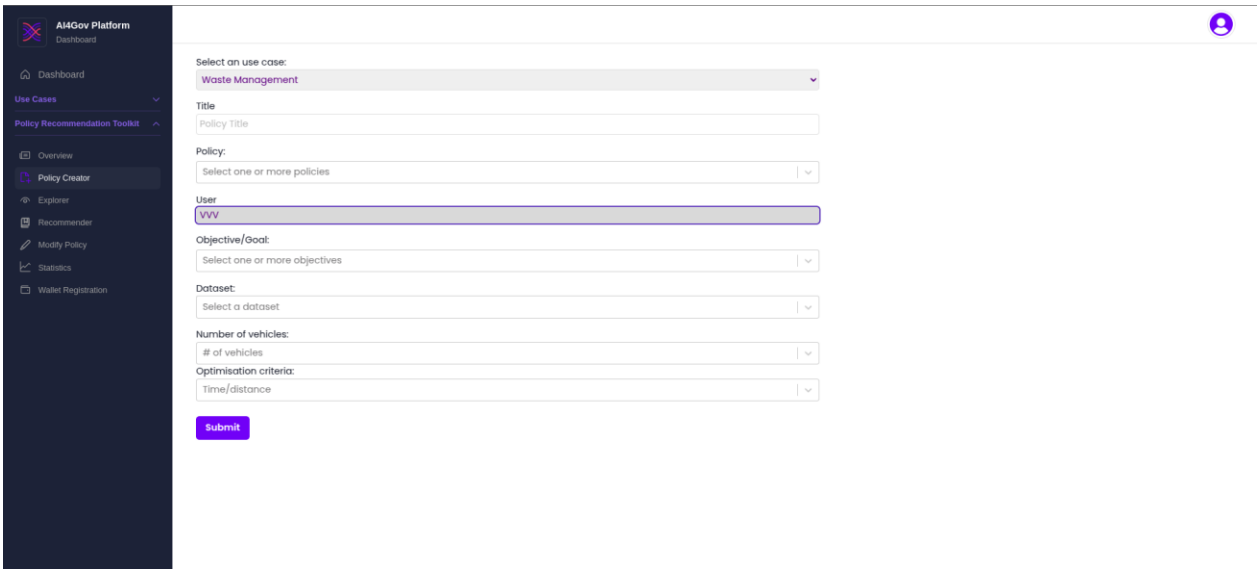

Figure 6: Policy Creator interface

The user can navigate among the existing policies (Figure 7). Then a policy can be chosen, and the user can examine and read the details of each policy and endorse the preferable policy (Figure 8).

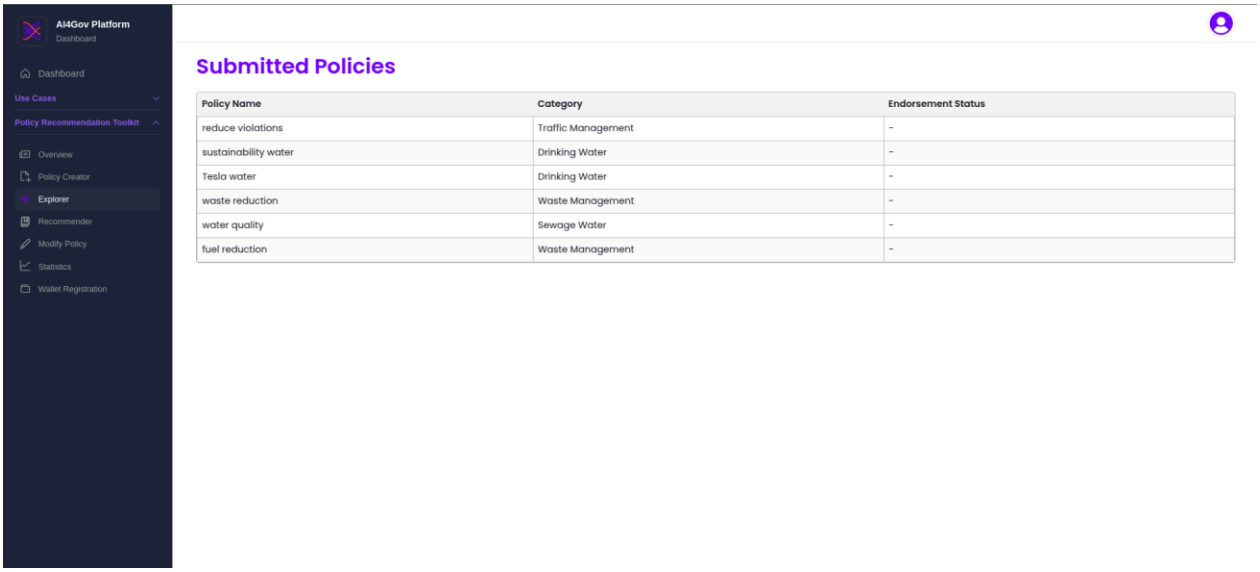

Figure 7: Policy Explorer

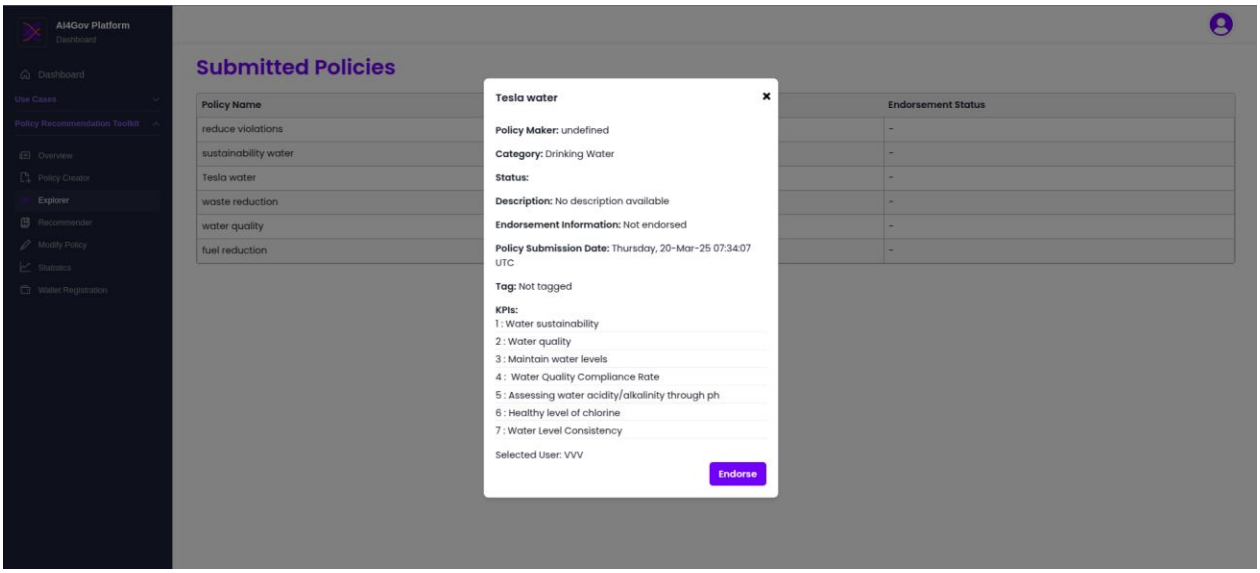

Figure 8: Policy Endorsement

The user is selecting KPIs as hard and soft constraints and retrieves as a response the policies that fulfil the KPIs of hard constraints, Figure 9. The retrieving policies have a rate regarding the soft KPIs that are fulfilled.

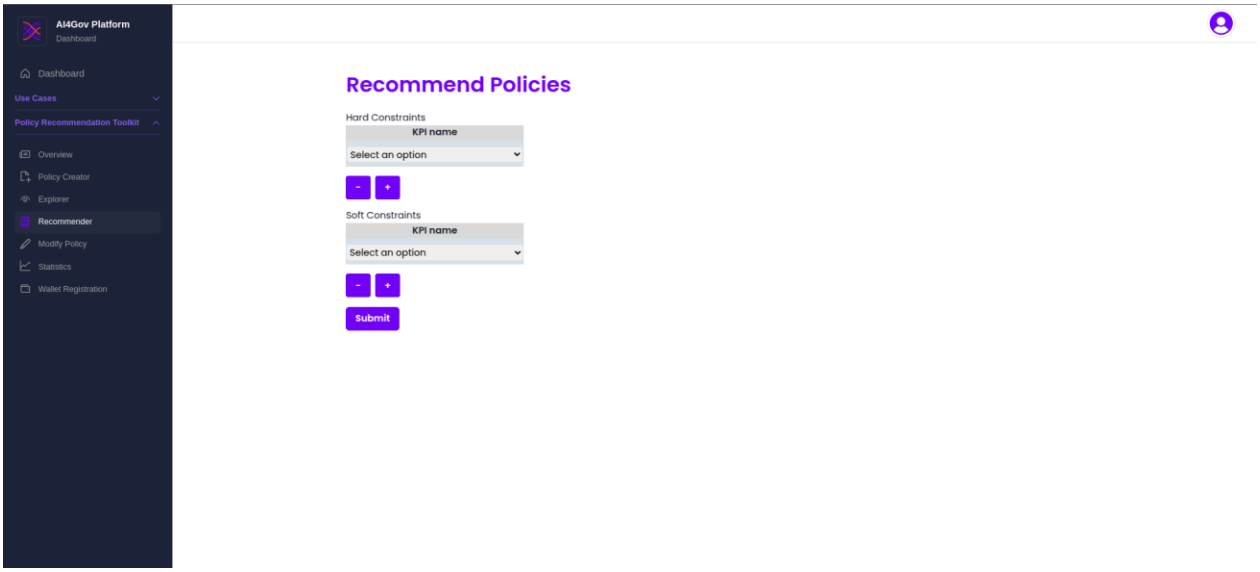

Figure 9: Recommendation results and KPI / constraints interface

There is an editing option via addition or re-definition of the relevant KPIs (Figure 10).

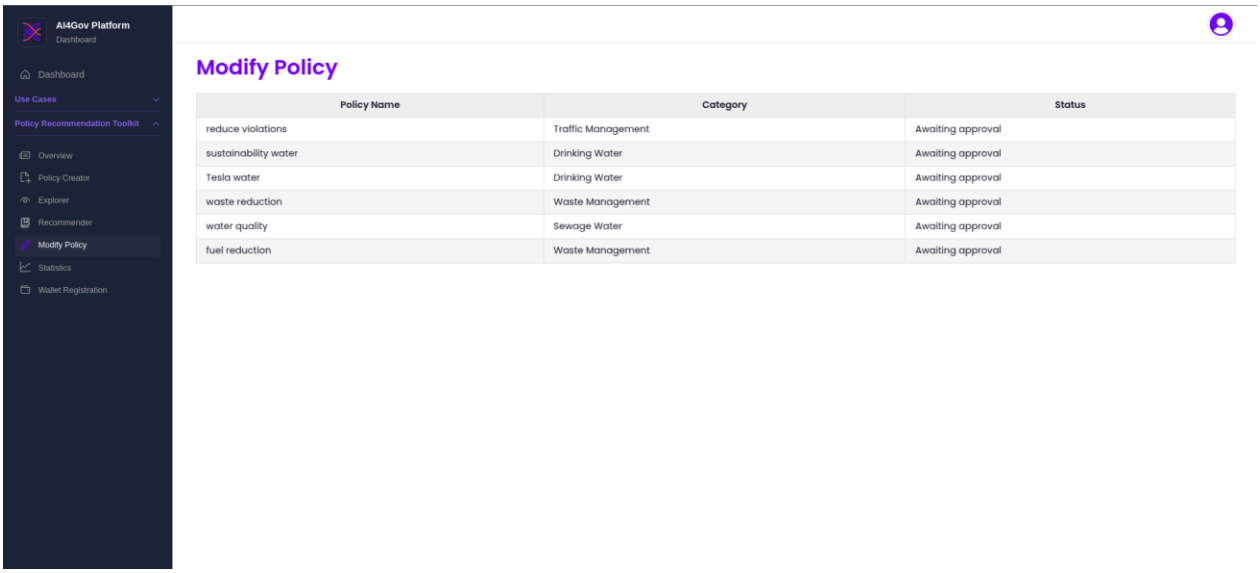

Figure 10: Modification of policies and editing interface

Also, there is a statistics menu that describes some analytics about the categories, the status and the endorsement of policies, Figure 11.

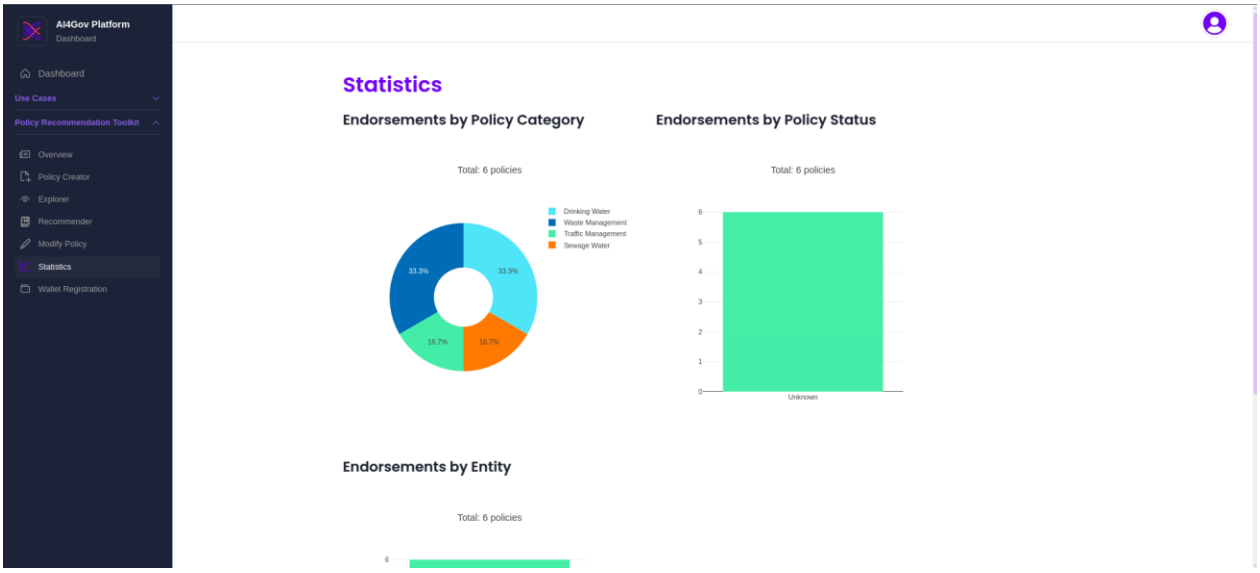

Figure 11: Comparative analytics

Finally, there is a menu about the credential issuance operation. The first step is to share a QR code for establishing a connection between the issuer (VVV or DPB) and the holder (citizen wallet), Figure 12. After the connection is established, the Issuer can fill the form about the attributes of holder/citizen and send the credential offer to the holder in order to store this to the wallet, Figure 13.

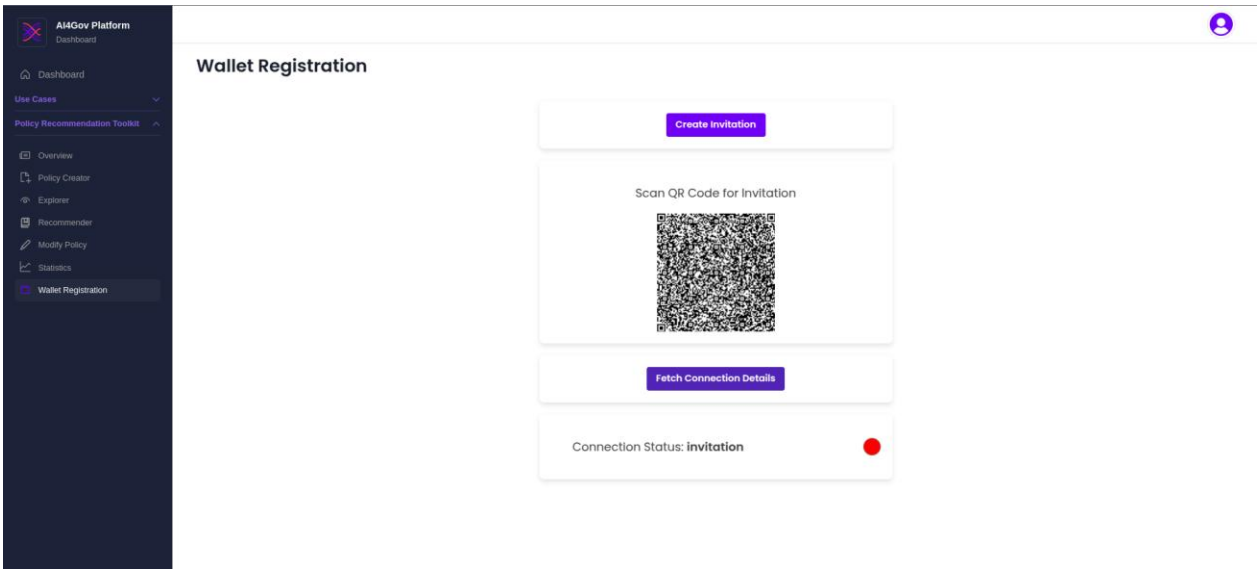

Figure 12: Sharing QR code for citizen invitation

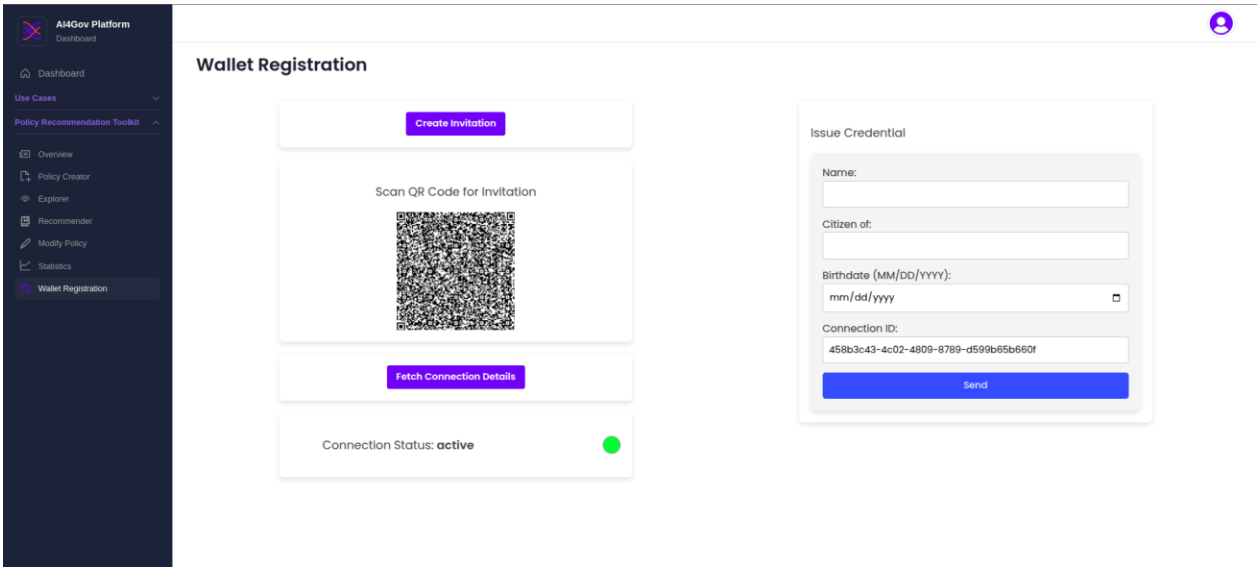

Figure 13: Form of attributes of verifiable credential

3.2.2 Objectives

The main objectives of the tool are the following:

- Facilitate policy definition and extension.
- Compare policies based on their performance on specified KPIs.
- Interface with the AI4Gov blockchain to execute smart contracts that execute commonly defined and endorsed business rules to evaluate existing or custom policies.

- Engaging Citizens via voting mechanism leveraging decentralized identity wallet with Verifiable Credential solution.

3.2.3 User Interaction

Figure 14 depicts the main use cases of the user interaction with the Policy Recommendation framework. The User can access the framework by logging in the visualisation workbench using the relevant functionality of the visualisation workbench and the authentication/authorisation schema that will be implemented via the AI4Gov access policy mechanism that will be deployed on the AI4Gov blockchain. The user can then access each of the framework's functionalities by choosing the relevant item from the workbench. These interactions involve viewing of policies, modification and extension of existing ones and custom policy definition and evaluation using the backbend's smart contract infrastructure.

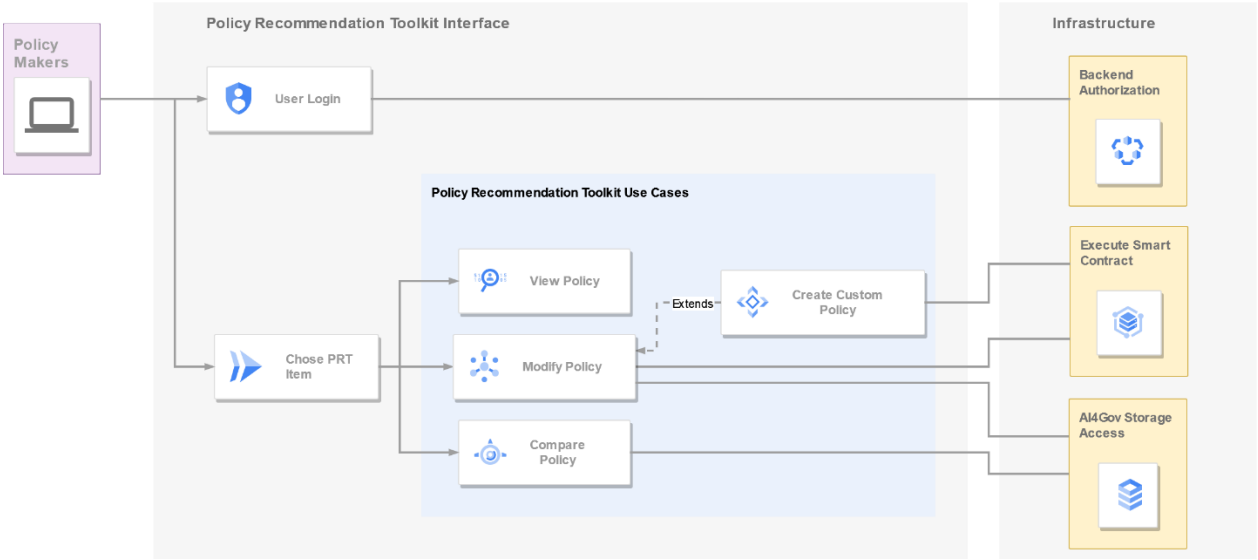

Figure 14: Use case diagrams for user interaction with the Policy recommendation toolkit

4 Bias Detector Toolkit Dashboard

4.1 Overview

The toolkit provides a holistic approach, combining a scrollytelling explanation of bias in AI, the catalogue (a visual synthesis of a variety of bias mitigation tools) and use-case specific solutions for data incompleteness in rare diseases. This tool aims to achieve 3 different goals, aiming to 3 different audiences: (1) increased awareness for general public; (2) provide bias-detection and mitigation tools for AI engineers; and (3) quantify bias in data for the patient-reported outcomes in rare diseases for rare disease policymakers and context experts.

4.2 Introductory Section

4.2.1 Overview

Introductory section is structured as a scrollytelling explaining AI, bias and bias in AI, followed by real life examples of the bias in AI.

4.2.2 Objectives

For the general public and policymakers, our strategy employs scrollytelling and real-life examples, to increase the understanding of the critical importance of bias mitigation.

4.2.3 User Interaction

The first section of the application is the scrollytelling part, where the user can be informed about bias and AI in general. Scrollytelling, also known as scroll-driven storytelling or scroll-based storytelling, is a web design technique that involves using the scrolling action of a webpage to reveal content in a narrative or visually engaging way. Instead of presenting information all at once on a single page, scrollytelling unfolds content gradually as the user scrolls down the page.

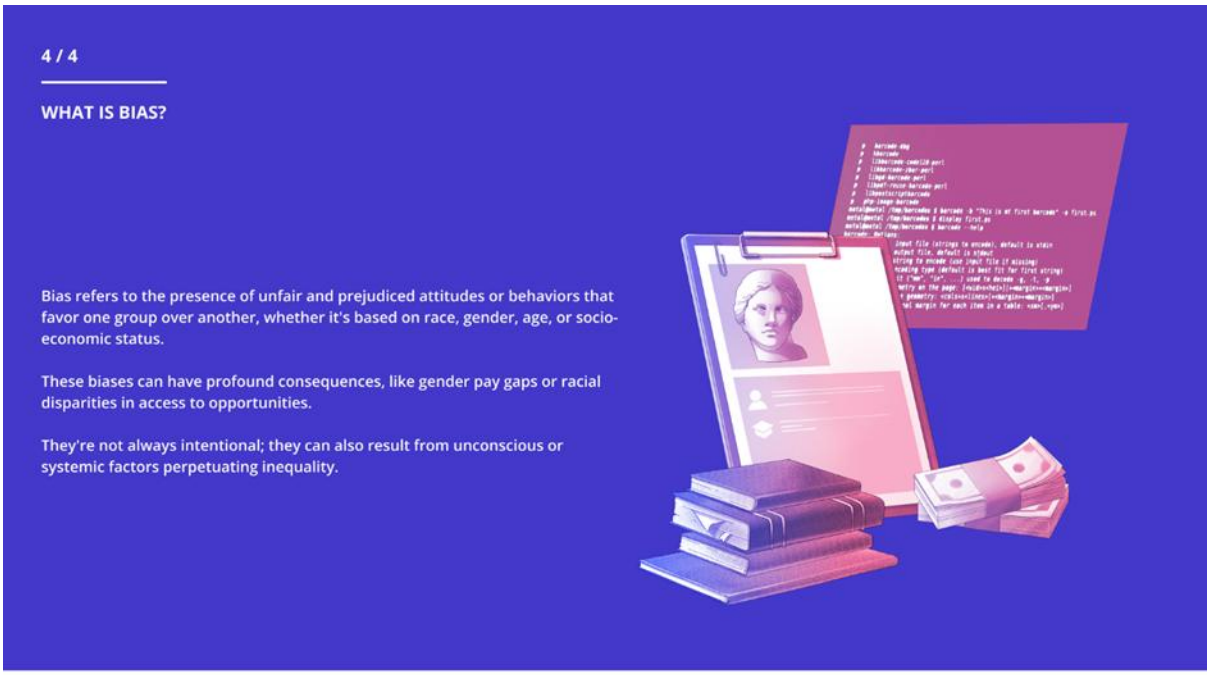

Figure 15: A frame from the scrollytelling part of bias

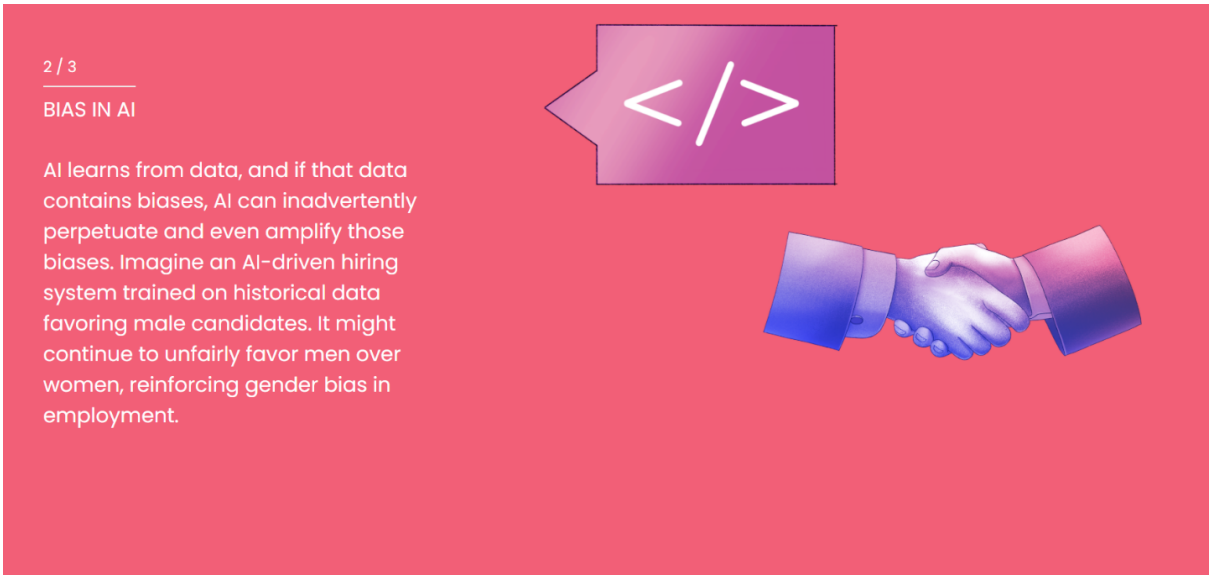

Figure 16: A frame from the scrollytelling part of bias in ai

The second section provides real examples of bias occurrences in different business sectors. Bias in real-world applications of ML has manifested in various forms, raising ethical concerns and highlighting the importance of responsible AI development. For each example, there is:

- a short description that summarizes the problem,
- the solution that either solved or mitigated the problem,
- reference material for further research on the topic.

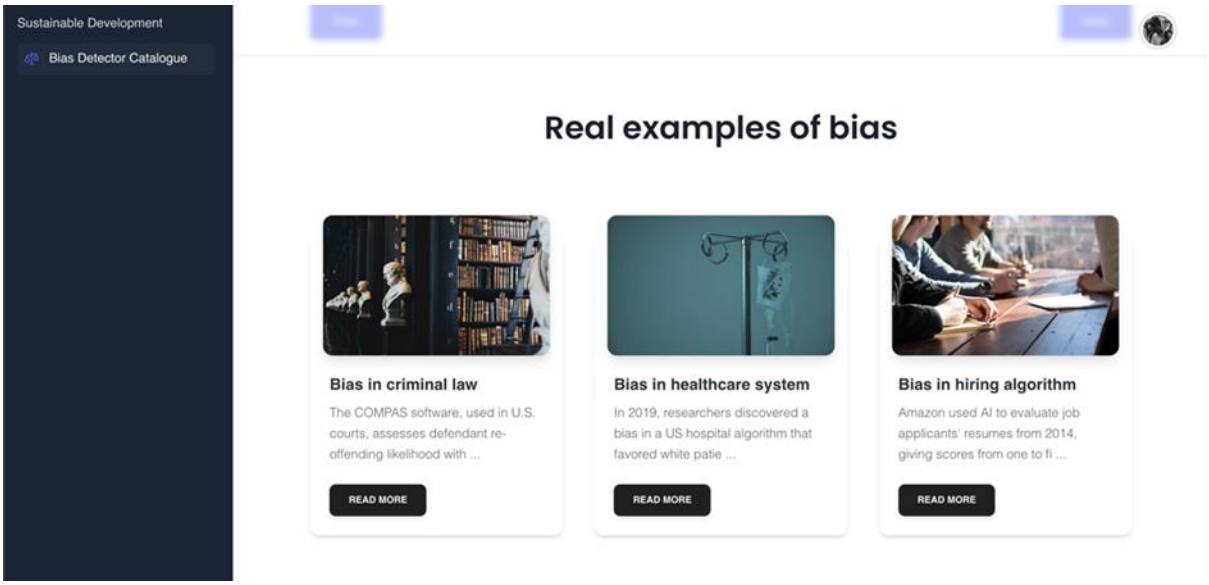

Figure 17: Real examples of bias occurences

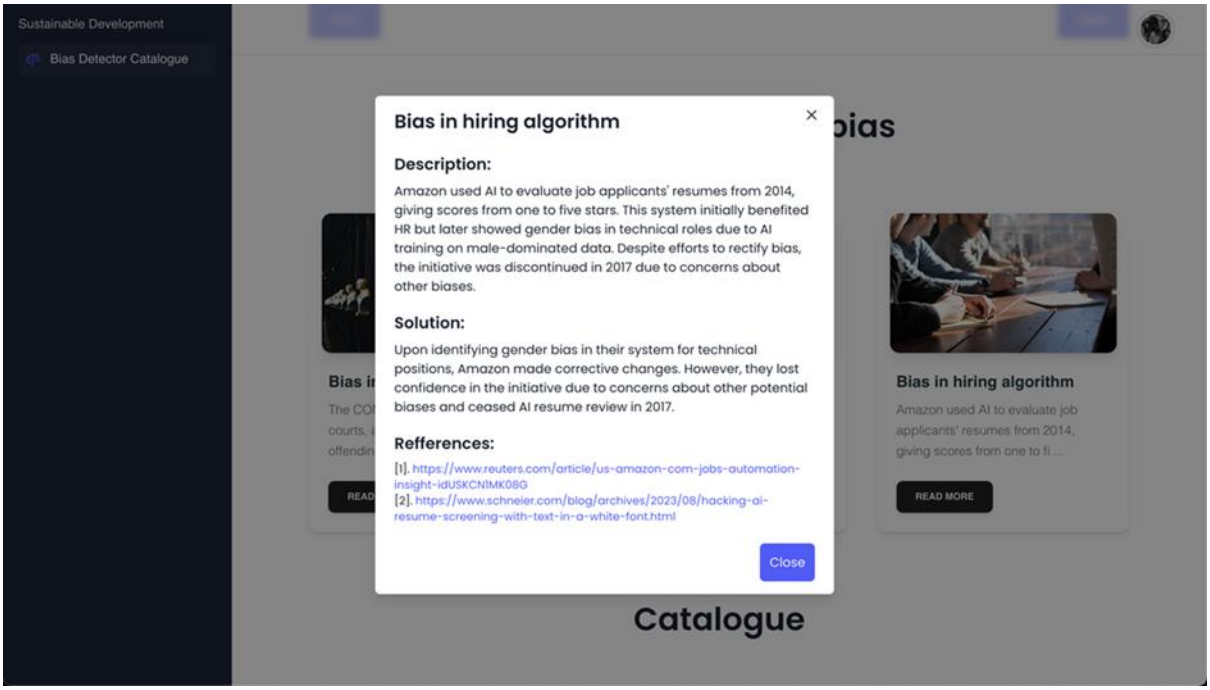

Figure 18: Details of each bias example

4.3 Bias Detector Catalogue Interface

4.3.1 Overview

The Bias Detector Catalogue employs visual synthesis to systematically analyse and present bias mitigation tools in a coherent manner. At its core, visual synthesis integrates data visualization principles, leveraging graphical representations to distil intricate information about various bias detection and mitigation tools. The catalogue utilizes visual elements such as graphs, matrices, and icons to categorize tools based on functionalities, applications, and performance metrics. Through this visual synthesis, users can efficiently navigate the diverse landscape of bias mitigation, gaining a nuanced understanding of each tool's strengths and limitations.

4.3.2 Objectives

For AI/ML engineers developing AI models, the objective is to present tools and methods for bias mitigation and detection as a catalogue to facilitate bias-free by design AI models.

4.3.3 User Interaction

Users are presented with each stage of training machine learning model's process. Upon clicking one, a list of available relevant projects is visualized. Every project has some properties such as Accuracy and Cost. The user can further expand each project's details by clicking on it and reveal additional information. This includes the project's source, type, programming language, a description, possible limitations and references.

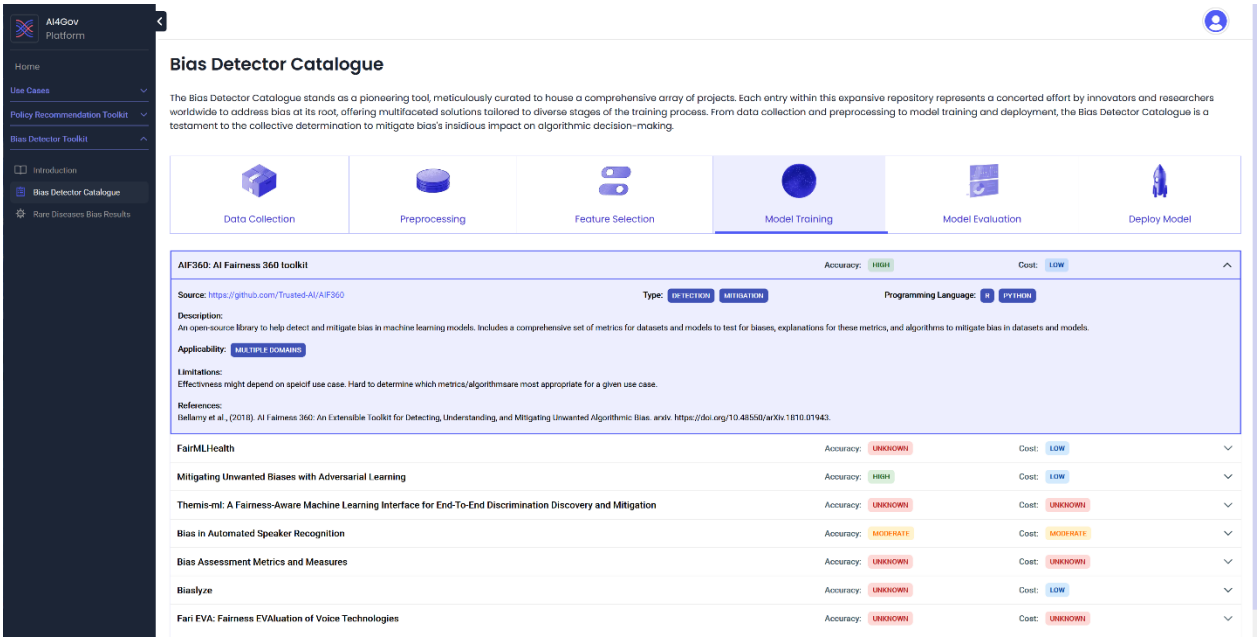

Figure 19: Bias Detector Catalogue Interface

4.4 Rare Diseases Bias Results

4.4.1 Overview

In rare diseases, limited scientific and clinical data often consist of single case reports and local data with small populations. Global cohorts' data is crucial, making patient-reported outcomes (PROs) and caregiver-reported outcomes (CROs) repositories essential for insights. It's hard to detect bias and data incompleteness in these repositories, and there are notable differences in data representation across continents. While a comprehensive understanding requires global data, the reality is that data predominantly comes from Western countries with more developed healthcare systems. This imbalance led our research team to develop an advanced analytics pipeline explicitly designed to quantify and address the incompleteness of data originating from regions with less developed healthcare systems.

4.4.2 Objectives

Objective is to develop methodology to systematically evaluate and eliminate or try to correct the bias inherent in PROs and CROs repositories.

4.4.3 User Interaction

Data incompleteness pipeline for rare diseases in available on the AI4Gov platform under the Rare Diseases Bias Results (Figure 19). User

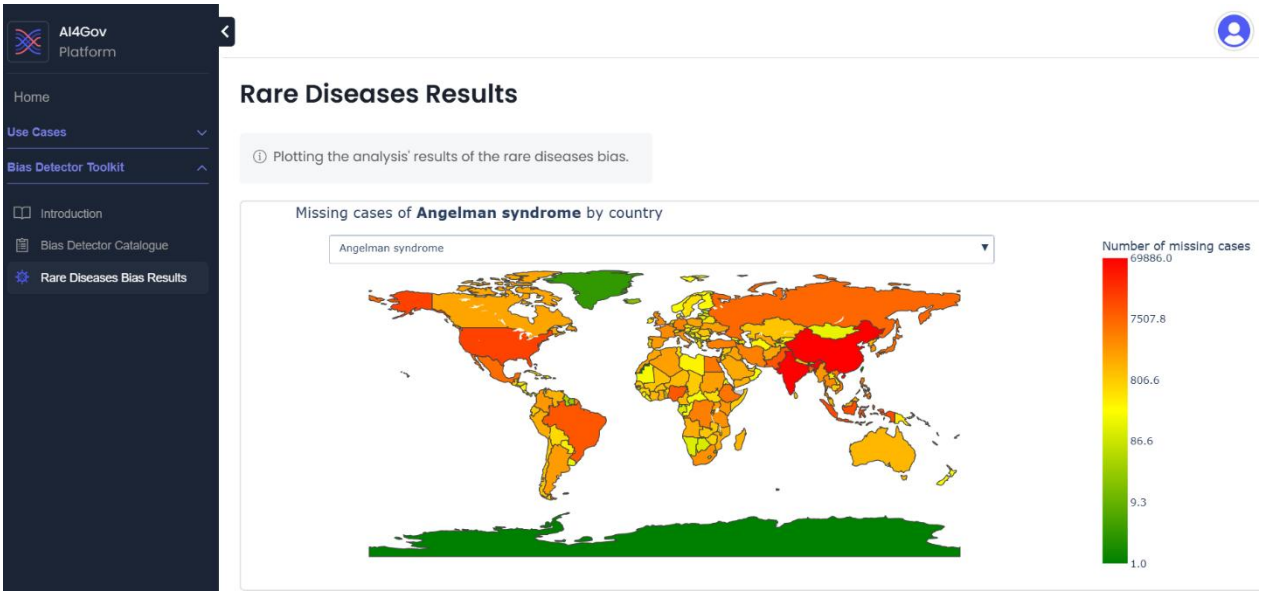

Figure 20: Interface for the rare diseases results

5 Water Management

5.1 Overview

The Badajoz Provincial Council's Environmental Consortium (DPB), acting as a vigilant entity focused on the ongoing supervision and effective management of water treatment processes, plays a pivotal role in addressing both wastewater and potable water concerns. In pursuit of enhanced efficiency and foresight, the consortium recognizes the significant advantages that a specialized tool for water quality monitoring and prediction can bring to its operations. This tool, designed with a specific focus on visualization, promises to revolutionize the way data is comprehended and utilized by various stakeholders within the consortium. From the diligent technicians operating in the intricate environments of Water Treatment labs to the administrative personnel at the heart of the Provincial Council, a dedicated visualization tool ensures seamless accessibility and interpretation of critical water quality information.

By facilitating a user-friendly interface that caters to the diverse needs of these professionals, the tool fosters a collaborative approach, streamlining communication and decision-making processes across the spectrum of roles involved in the management of water resources. In essence, this innovative solution emerges as an indispensable asset, harmonizing the efforts of the consortium's workforce for the comprehensive and sustainable treatment of water resources in the Badajoz region.

5.2 Use Cases

5.2.1 Drinking Water Forecasting

For this use case, a tool is being developed to assist technicians in identifying potential sources of inefficiency within the system related to water quality. The tool will be implemented through a user interface designed to facilitate the input and analysis of data from drinking water treatment stations. Using data on parameters such as pH, chlorine levels, water level, and the quantity of instantaneous water output, the system will generate a 24-hour prediction for these variables.

This prediction will be accompanied by a heat map and an explanation of how to interpret it. The generated chart will help users understand the relationships and influences between the variables that contribute to the prediction.

The resulting chart provides users with valuable insights into the expected behaviour of selected drinking water parameters, supporting proactive decision-making for water quality management and resource planning. The user-friendly interface, combined with advanced forecasting capabilities, makes it a powerful tool for informed decision-making, leveraging historical trends and predictive analytics for better management of drinking water quality.

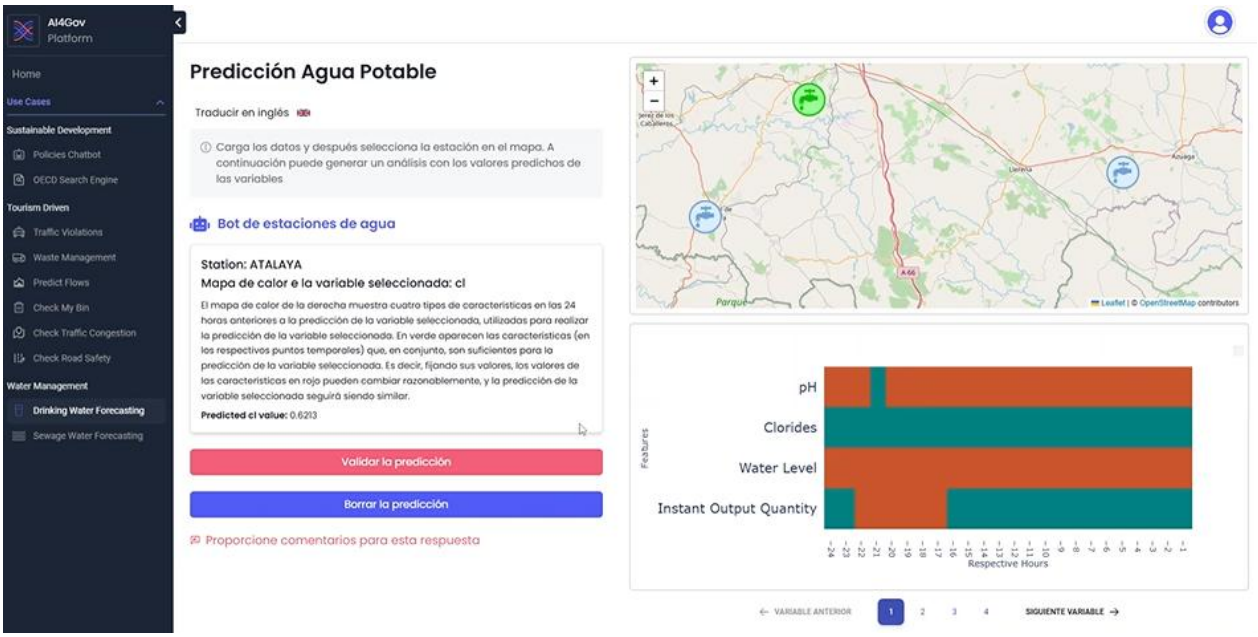

Figure 21: Drinking Water Interface

5.2.2 Sewage Water Forecasting

For this use case, a user interface has been developed that provides a robust platform for entering and analysing wastewater variables, with a focus on critical parameters. At the same time, users can input specific energy-related variables associated with wastewater treatment plants. By analysing historical data from these plants, the use case offers technicians a tool to identify inefficiencies in energy consumption.

Once the data is submitted, the analytics services employ advanced time-series forecasting algorithms to generate a predictive graph along with a detailed explanation of the relationship and significance of the variables entered into the prediction. This graph offers a practical seven-day forecast, enabling proactive decision-making in wastewater treatment. The tool's intuitive interface, combined with its advanced forecasting capabilities, positions it as an invaluable resource for making informed decisions regarding wastewater treatment plants.

Ultimately, the pilot program helps policymakers identify recurring issues and trends, providing valuable insights to guide long-term investment strategies in water cycle management.

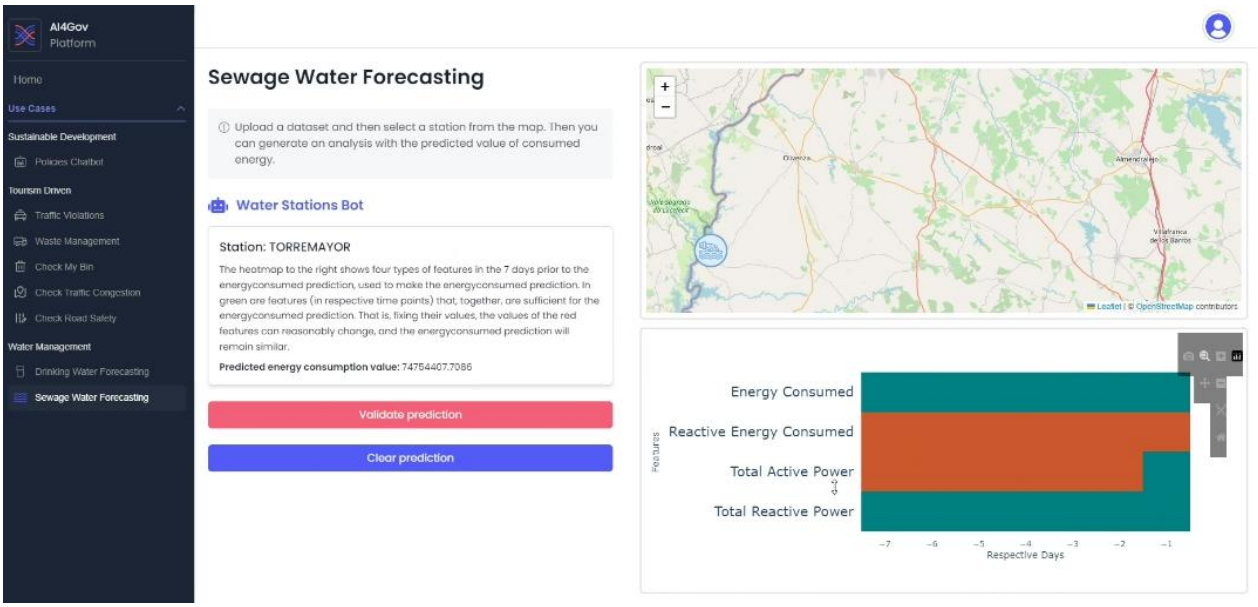

Figure 22: Sewage Water Interface

6 Sustainable Development

6.1 Overview

International Research Centre in Artificial Intelligence (IRCAI), under the wing of JSI, is an organization under the auspices of UNESCO that aims to promote international cooperation and collaboration in the development and deployment of AI for the benefit of humanity. Its vision is to bring AI solutions to Sustainable Development Goals, and under AI4GOV the plan is to develop both methodology for better bias identification and mitigation in the Top100 projects use case, and for development of SDG observatories.

6.2 Use Cases

6.2.1 Top 100 projects checklist for applicants and reviewers

Top100 projects is an IRCAI initiative, funding projects from all over the world to address problems related to the SDGs by using AI. The primary objective is to identify and present solutions worldwide, making a significant contribution to the SDGs through the creation of a vast platform for sustainable solutions. To introduce a methodology for the evaluation of bias in project proposals and how the applicants address these issues, direct output of the Bias Detector Catalogue are checklists both for applicants and reviewers. Applicants to the Top100 projects are asked to fill out an online form in the process of their application. Reviewers, which are the members of the IRCAI Scientific Programme Committees, Scientific Journal Editorial Board, and Business and Impact Council, are similarly presented with a checklist upon review of application, that enables them to evaluate to what extend bias component has been considered in the models of the Top100 applications. Both checklists have been implemented as additional questions in the previous-existing application and revision forms.

The Bias detector catalogue, presented in Chapter 4, is also being used by this use case, for the developers of Top100 projects to gain insights on the tools they can use.

6.2.2 SDG Observatories Methodology Interface

IRCAI SDG observatories are a tool that monitors achievements of SDGs, more specifically how AI is making contribution towards reaching these goals. The tool is under development, and methodology to detect bias and mitigate it at data level have been developed for the rare disease observatory. Contribution of visualization workbench to this use case has been showcasing the results of the data incompleteness pipeline for rare diseases, presented in Chapter 4.

Part of this use case is also the alcohol in traffic exploration of bias in data, that has been presented in D6.3 and will be further presented in D6.4.

6.2.3 Citizen Engagement Using NLP Interface

Policies Chatbot Interface is a platform designed to facilitate comprehensive insights into AI advancements across various countries, based on the AI policy documents. Based on the first validation phase, the interface has been refactored (Figure 22) and renamed to OECD Search Engine. Refactoring improved speed and added the functionality of showing the whole pdf document alongside with the possible answers (Figure 23).

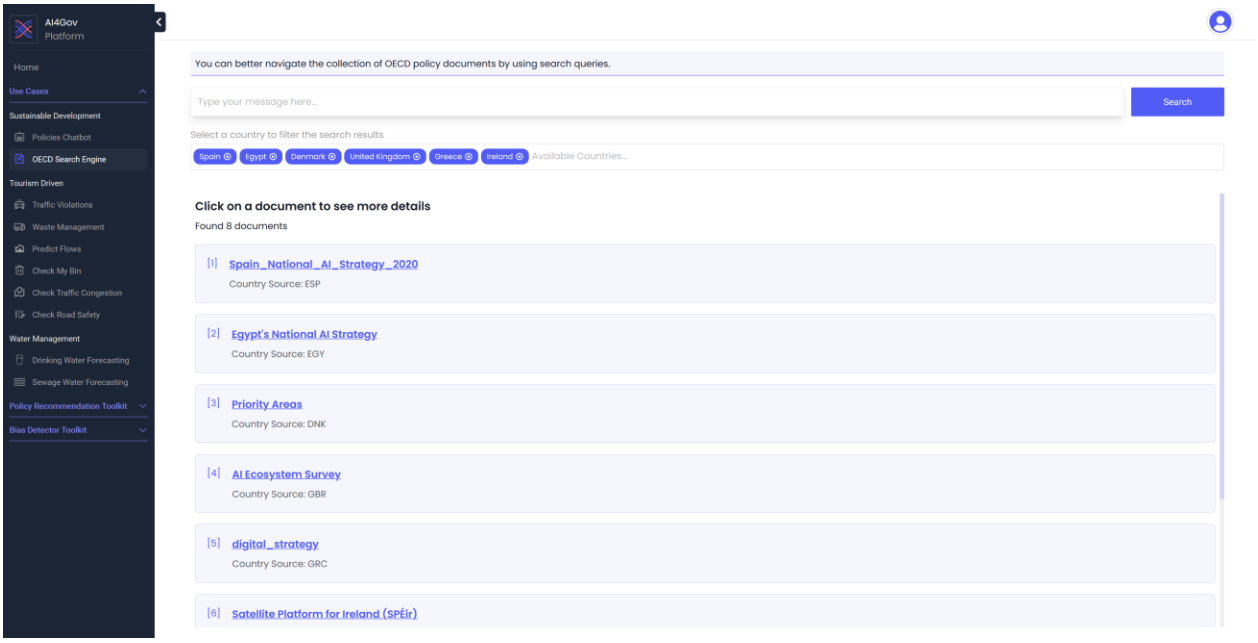

Figure 23: OECD Search Engine Interface

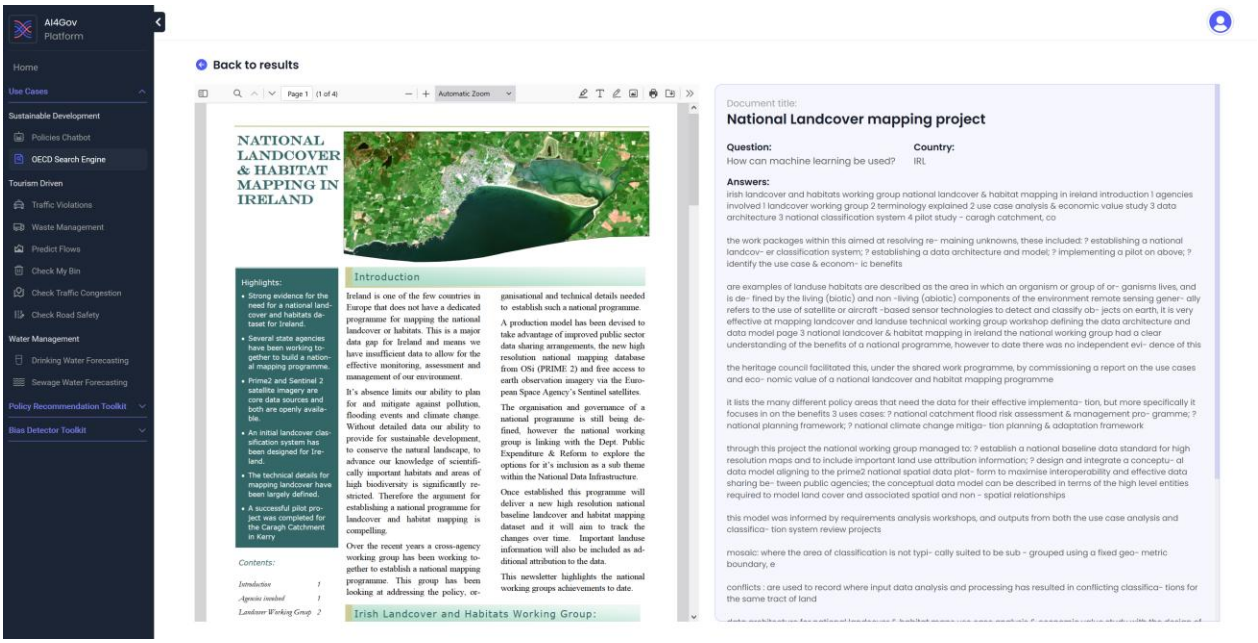

Figure 24: OECD Search Engine showcasing document

7 Tourism-driven Policy Management and Optimisation

7.1 Overview

The Municipality of Vari-Voula-Vouliagmeni (VVV) is a Local Authority with approximately 600 personnel that is controlled by the Municipal Council and the Mayor. The municipality's thriving tourist appeal necessitates a thorough revamp and enhancement of its policies based on the flow of visitors and tourists, taking into consideration both spatial and temporal factors (such as specific locations at particular times) and responding to specific scenarios as they arise. The Ministry of Tourism will serve as an additional source of information, supplying data on seasonal visitors and tourists, offering insights into the expected influx and the corresponding temporal dynamics of their visits.

In both of its use cases the Interactive visualization workbench will be an extremely supportive tool as it will provide an estimation of effects of policies, while delivering a user-friendly visualization on interactive analytics based on the data generated. Moreover, the workbench will prompt the users to draw their own predictions through trinity visualization by adjusting their prior knowledge. Finally, it will automatically recommend the most representative visualizations based on self-explanation capabilities.

7.2 Use Cases

7.2.1 Traffic Violations Interface

The Municipality has to deal with traffic violations especially those related to illegal Parking that show a significant increase during the summer and weekends due mainly to the flow of visitors and tourists. The Municipal administration has developed a Sustainable Urban Mobility Plan to improve the overall quality of life for residents and visitors by addressing major challenges related to congestion, road safety and parking. To this end, a tool has been developed within the framework of the project that monitors the number of traffic tickets issued and analyses their time and spatial evolution in order to allocate municipal police staff, vehicles and equipment in an optimum way.

This user interface serves as a comprehensive platform designed for efficient predictions of potential violations within the VVV municipality. The user is presented with a user-friendly form that facilitates the input of key parameters crucial for generating accurate predictions. The form comprises four essential fields: month, allowing the user to specify the temporal aspect of the potential violation; weekday or weekend, aiding in further narrowing down the time frame; time of day, enabling the user to specify whether the violation is likely to occur in the morning, noon, evening, or night; and violation type, offering a selection of violation categories such as speed limit violation, traffic light violation, parking violations, and more.

Upon submission of this detailed information, the analytics services employ a sophisticated algorithm (in this implementation Decision Tree) to analyze historical data and patterns specific to the municipality. Leveraging machine learning and predictive modeling, the interface generates

a dynamic map that visually highlights the area within the municipality premises most susceptible to the specified violation based on the provided parameters. The prediction is presented in a visually intuitive manner, with color-coded or shaded region the area.

Users can leverage this predictive map to enhance proactive measures, such as increased surveillance or targeted awareness campaigns in the identified high-risk zones. The interface thus not only streamlines the reporting process but also empowers local authorities and residents with valuable insights to foster a safer and more compliant community. This innovative approach to predictive mapping in the context of municipal violations demonstrates the integration of technology and data analytics for the betterment of public safety and urban planning.

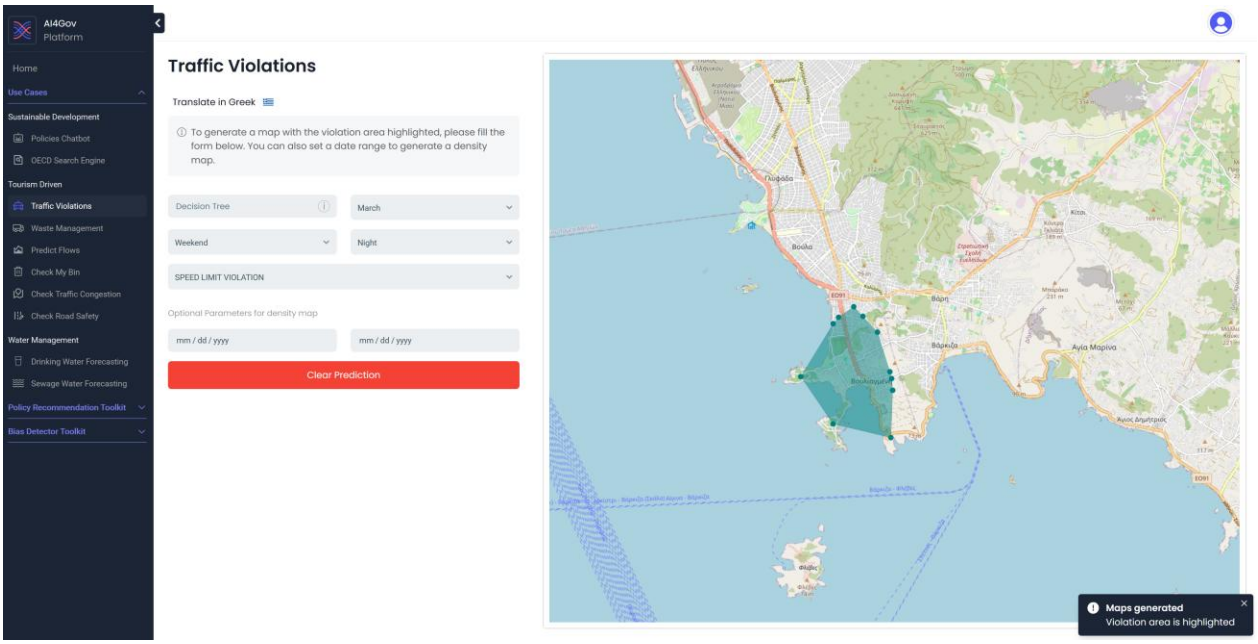

Figure 25: Traffic Violations Interface 1

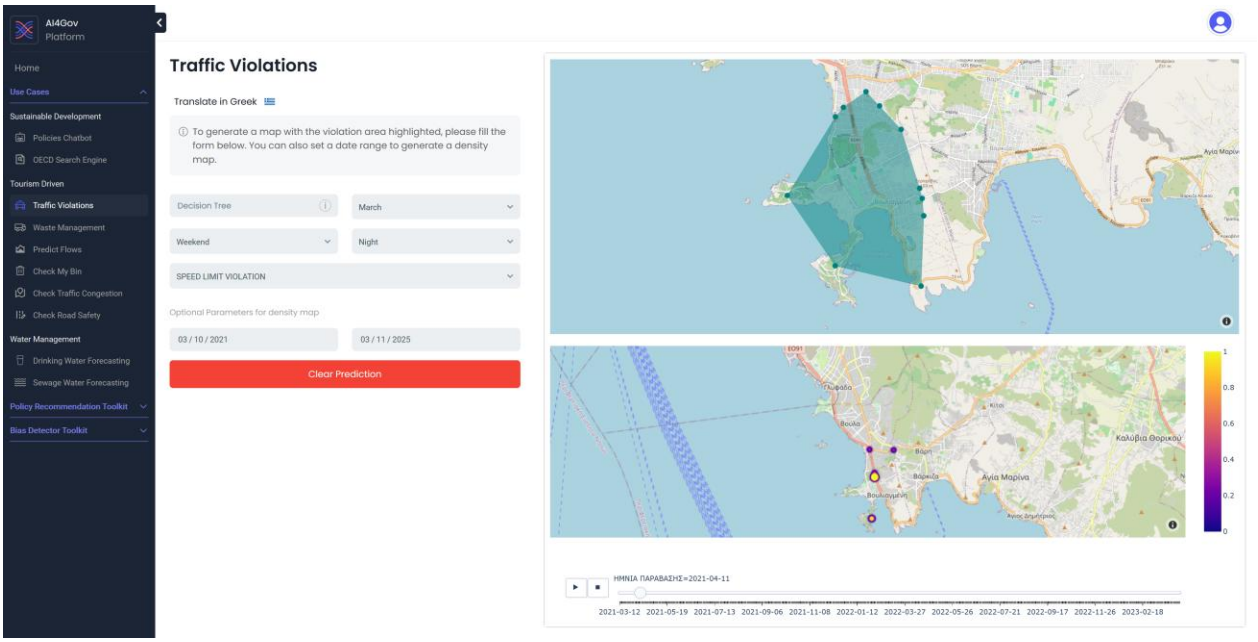

Figure 26: Traffic Violations Interface 2

7.2.2 Waste Management Interface

As the Municipality of Vari-Voula-Vouliagmeni faces larger amounts of waste during the summer and weekends, due mainly to the flow of visitors and tourists, it requires an innovative tool to streamline waste monitoring and optimize the allocation of waste management staff and resources, while implementing the Pay As You Throw System in an efficient way.

This waste management interface for the VVV project pilot presents a comprehensive solution for optimizing the collection of different types of waste. The user is greeted with an intuitive interface that allows them to select the type of waste bin they intend to service - whether it'll be Brown for organic waste, or Green for general waste. Upon the selection of the waste type, the system dynamically generates an optimal path for the fleet of waste trucks to follow, minimizing time and fuel consumption. The interface intelligently considers factors such as bin proximity and historical waste generation patterns to generate an efficient collection route.

As the system generates the route, it simultaneously compiles a list of bin IDs in the order they should be serviced. Additionally, the interface provides real-time metrics, including the estimated amount of waste (in litters) that will be collected during the route. This feature aids in resource planning and ensures that the waste trucks are appropriately equipped for the volume they will encounter. Moreover, the total time required for the waste collection process is calculated and displayed, allowing for better scheduling and optimization of the overall waste management operations.

Throughout the entire process, a live map is displayed on the side of the interface, providing a visual representation of the bin locations within the Vari - Voula – Vouliagmeni municipality. This visual aid enhances user understanding and allows for quick assessment of the proposed collection route. The combination of route optimization, real-time metrics, and a dynamic map

creates an integrated and user-friendly waste management interface that significantly contributes to the efficiency and sustainability of the municipality's waste collection operations.

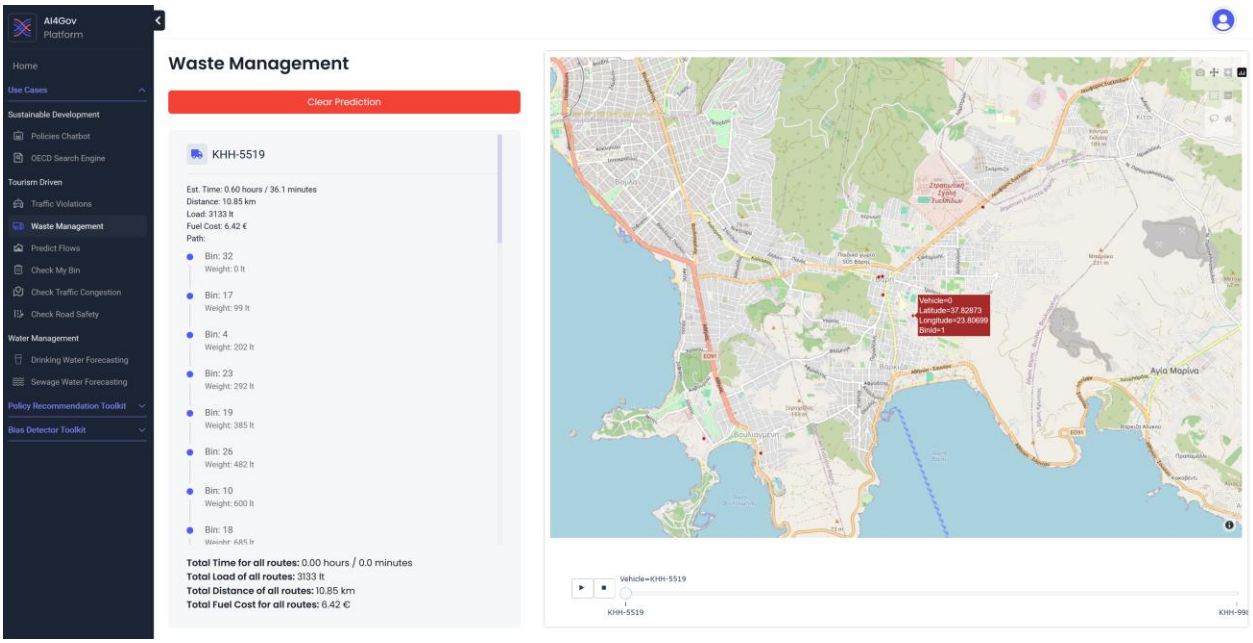

Figure 27: Waste Management Interface

7.2.3 Predict Flows Interface

The fill level of the Municipality's bins can also be utilized in the context of touristic policies, since the rate at which a bin fills is highly co-related with the flows of the citizens/tourists in the area where this bin is located. In other words, the smart bins that are located in higher populated areas and areas that are visited by more people fill at a higher rate than the smart bins that are located in less populated areas.

To this end, a new interface was designed for predicting the touristic flows in the premises of VVV. More specifically, a neural network is trained to provide predictions for the future fill level of the smart bins. Then the user can choose to split the municipality in a selected number of areas, based on the location of the bins and also select a number of clusters based on the predicted fill levels and the corresponding fill rate. After that, a map is generated that highlights with different colors the clustered areas in terms of the projected fill rate, thus allowing the policy makers to identify expected number of tourists per area of interest and create/enhance the corresponding policies.

In conclusion, predicting the touristic flows in an area is of vital importance for the formulation of the corresponding policies that will enhance both the experience of the visitors and ease the everyday lives of the locals.

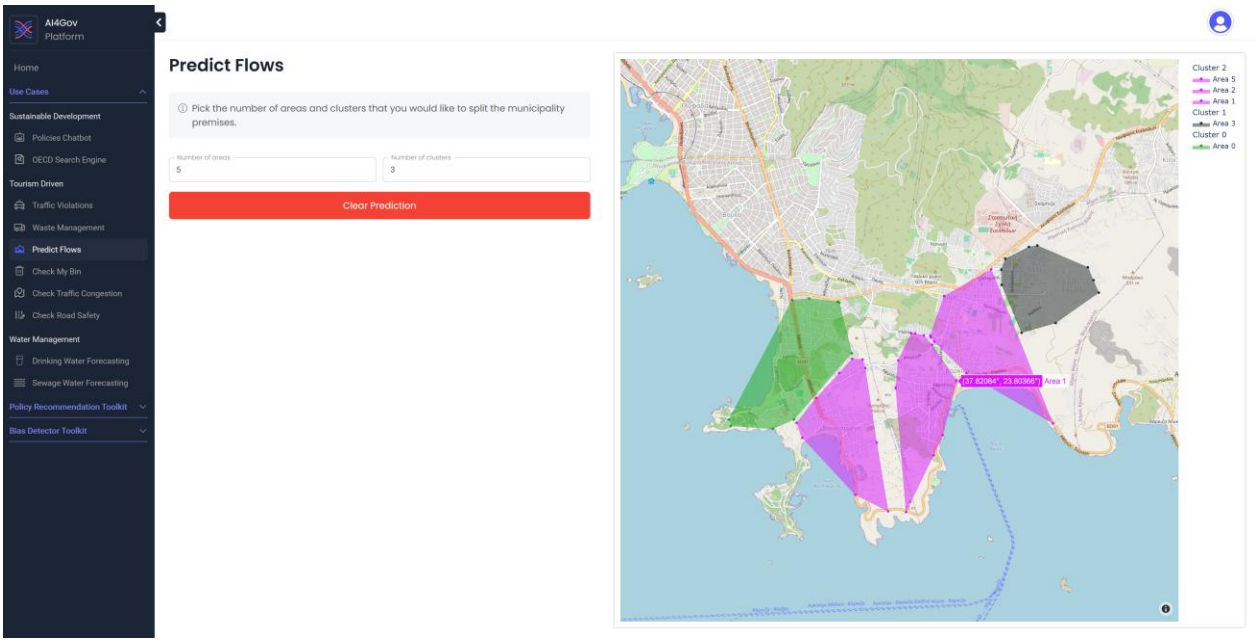

Figure 28: Predict Flows Interface

7.2.4 Check My Bin Interface

Within the Check My Bin interactive interface, designed specifically for the citizens and businesses of the Vari - Voula – Vouliagmeni municipality, users can easily monitor and manage waste disposal more efficiently. By selecting any of the municipality's smart bins, individuals gain access to crucial real-time data regarding its fill level and anticipated collection schedule.

Upon selection, a detailed prompt appears, offering insights into whether the bin is scheduled for emptying in the next collection round or if it remains within acceptable fill thresholds and will not be serviced yet. This prediction is powered by a neural network trained on the bin’s historical data, enabling the system to provide highly accurate forecasts based on past usage patterns, seasonal trends, and external factors such as tourism activity. This predictive functionality is particularly beneficial for both residents and businesses, especially those in the tourism sector, as it enables them to plan waste disposal proactively.

By knowing exactly when a bin will be emptied, users can avoid unnecessary overflow, contribute to a cleaner urban environment, and optimize their waste management practices in alignment with municipal services. This data-driven approach not only enhances convenience but also promotes sustainability by reducing inefficiencies in waste collection and ensuring smarter resource allocation within the city.

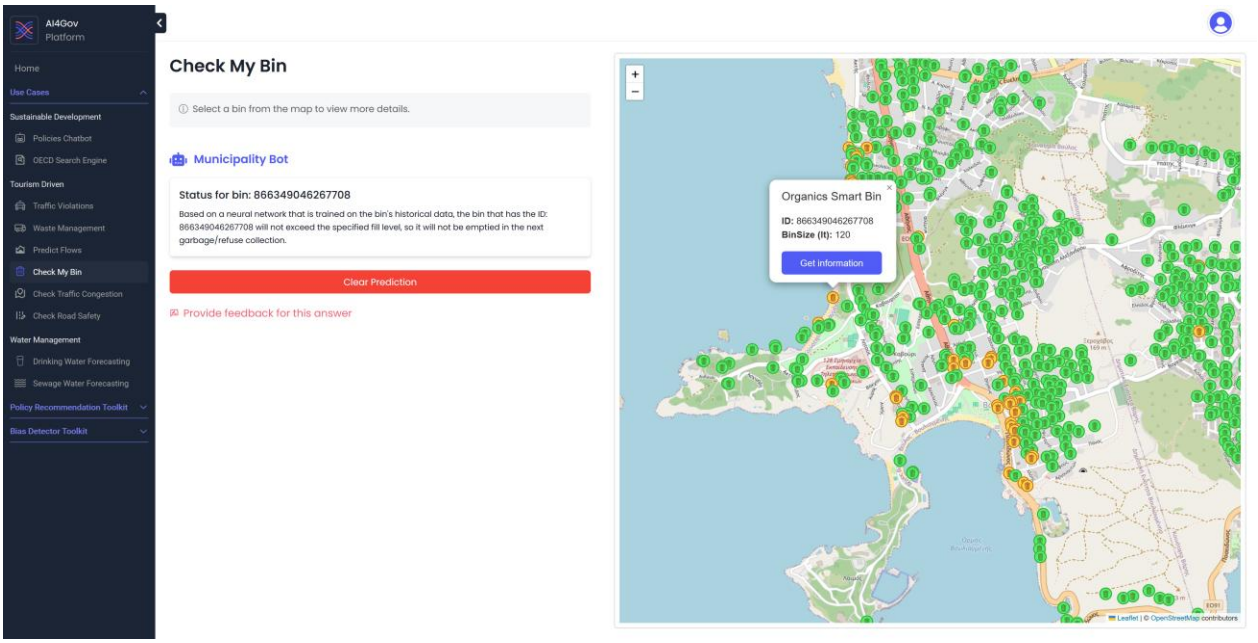

Figure 29: Check My Bin Interface

7.2.5 Check Traffic Congestion Interface

Within the Check Traffic Congestion interface, citizens and visitors of the Vari - Voula – Vouliagmeni municipality can efficiently plan their mobility and parking needs with the help of advanced AI-driven predictions.

By selecting a specific date and time, users can receive real-time insights into expected traffic conditions and parking availability within the municipality premises. Once the user inputs their preferred time, a powerful neural network—trained on historical traffic and parking data—analyzes patterns, past trends, and external factors such as seasonal fluctuations to generate accurate forecasts. The results are then displayed on an interactive map, highlighting the most suitable roads for parking at the given time while also marking areas to avoid due to high congestion or limited availability.

This predictive system empowers users to make informed decisions, reducing unnecessary time spent searching for parking and easing overall traffic congestion in the town. By optimizing mobility and transportation efficiency, the interface not only enhances user convenience but also contributes to a more sustainable urban environment by minimizing emissions and improving traffic flow within the municipality.

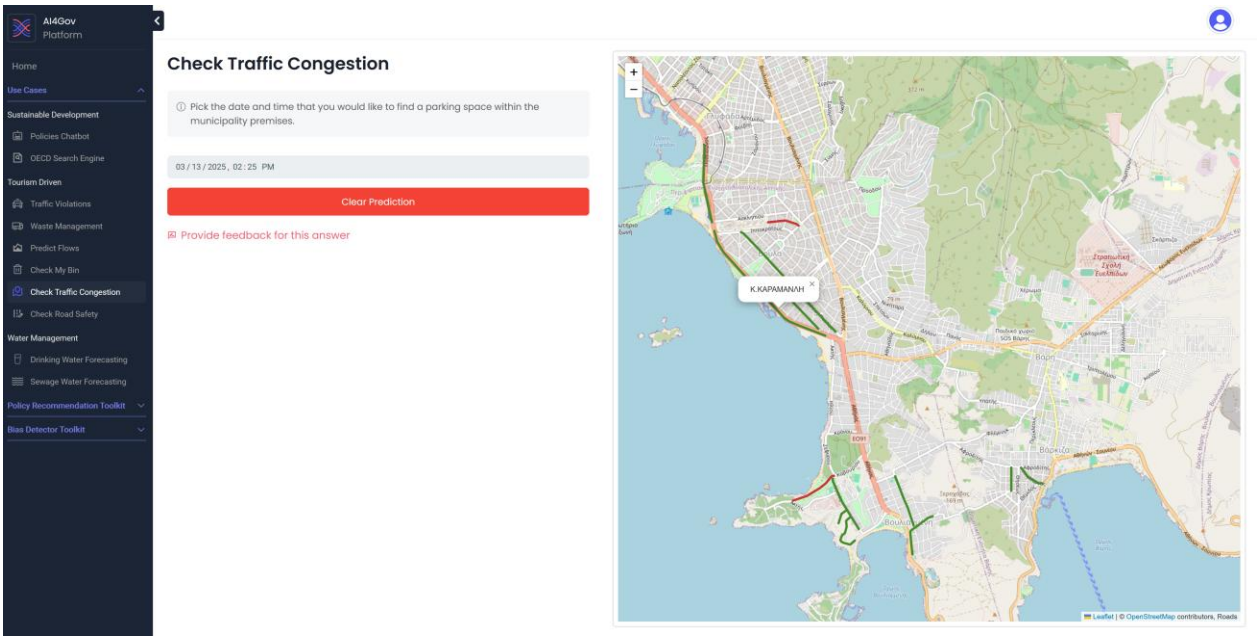

Figure 30: Check Traffic Congestion Interface

7.2.6 Check Road Safety Interface

Within the Check Road Safety interactive interface, citizens and visitors of the Vari - Voula – Vouliagmeni municipality can access crucial insights into road safety conditions, helping them make informed decisions about their routes.

By selecting a specific date and time, users receive predictions about the safety levels of different roads within the municipality. A sophisticated neural network, trained on historical traffic patterns and other relevant data, analyzes potential risks and generates an interactive map. This map visually highlights roads that are considered safe and recommended for travel, while also marking dangerous areas that should be avoided due to factors such as high accident probability, poor lighting, or adverse weather conditions.

By providing this valuable information, the system enables users to adapt their travel plans proactively, reducing the likelihood of accidents and improving overall road safety. Whether for daily commuting or tourism purposes, this AI-driven approach enhances public awareness, promotes safer driving behavior, and contributes to a more secure and efficient transportation network within the municipality.

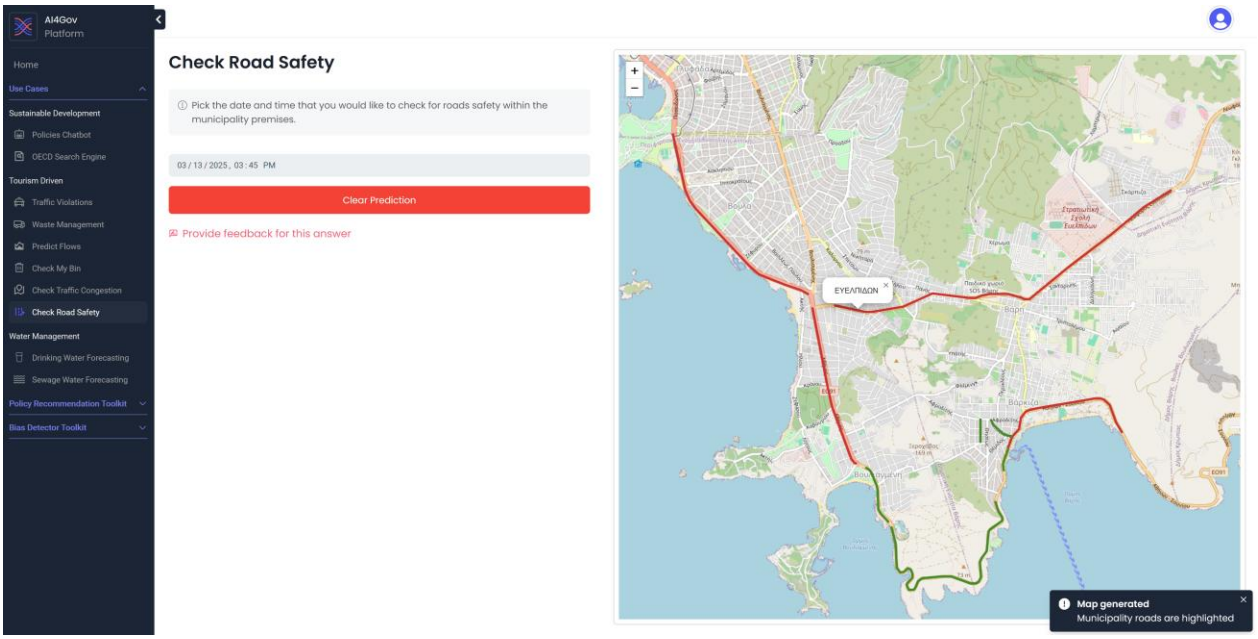

Figure 31: Check Road Safety Interface

8 Conclusions

In conclusion, the updated Visualization Workbench represents a significant leap forward in harnessing the transformative potential of AI across diverse sectors. As the enhanced front end of the AI4Gov project, it continues to serve as the central gateway to an expanding suite of use cases in Sustainable Development, Tourism, and Water Management, reflecting the collaborative efforts of visionary partners. With new advanced functionalities, the Visualization Workbench is more accessible, adaptable, and powerful than ever before.

Beyond merely hosting use cases, the Workbench amplifies their impact through its upgraded core services—the Bias Detector Toolkit and the Policy Recommendation Toolkit—which now offer deeper analytical insights and enhanced visualization layers. These improvements further solidify its role in promoting ethical AI, transparency, and informed decision-making within governance.

9 References

Docker. (n.d.). Retrieved from https://www.docker.com/.

Kubernetes. (n.d.). Retrieved from https://kubernetes.io/.

Law, P. M. (2020). Designing Tools for Semi-Automated Detection of Machine Learning Biases: An Interview Study. Retrieved from https://doi.org/10.31219/osf.io/uvjqh

leikas, J. J. (2022). Governing Ethical AI Transformation: A Case Study of AuroraAI. *Frontiers in Artificial Intelligence*. Retrieved from https://doi.org/10.3389/frai.2022.836557

Ltifi, H. A. (2016). Fuzzy logic–based evaluation of visualizations generated by intelligent decision support systems. *Information Visualization*, 17(1), 3-21. Retrieved from https://doi.org/10.1177/1473871616674046

Mikhaylov, S. E. (2018). Artificial intelligence for the public sector: opportunities and challenges of cross-sector collaboration. *Philosophical Transactions of the Royal Society A: Mathematical, Physical Engineering Sciences*. Retrieved from https://doi.org/10.1098/rsta.2017.0357

Next.js. (n.d.). Retrieved from https://nextjs.org/.

Wu, X. P.-M. (2020). An interactive web-based geovisual analytics platform for co-clustering spatio-temporal data. *Computers & Geosciences*, 137. Retrieved from https://doi.org/10.1016/j.cageo.2020.104420