

AI4Gov

Trusted AI for Transparent Public Governance
fostering Democratic Values

Deliverable 4.3

Policies Visualization Services V1


28-12-2023

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	 <p>This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0). See: https://creativecommons.org/licenses/by-nc/4.0/</p>

AUTHORS		
	Name	Organisation
Document leader	Andreas Karabetian	UPRC
Participants	Miguel López Corbacho	DPB
	Dimitris Apostolopoulos	VVV
	Alenka Guček	JSI
	Dimitris Ntalaperas	UBI
Reviewers	Theodore Chadjipadelis	AUTH
	Georgia Panagiotidou	AUTH
	Vicky Stroumpou	MT
	Stavroula Kefala	MT

VERSION HISTORY				
Version	Date	Author	Organisation	Description
0.1	27/10/23	Andreas Karabetian	UPRC	Initial ToC
0.2	02/11/23	Andreas Karabetian	UPRC	Updated ToC
0.3	07/11/23	Miguel López Corbacho	DPB	Updated Section 5
0.4	16/11/23	Dimitris Apostolopoulos	VVV	Updated Section 7
0.5	30/11/23	Alenka Guček	JSI	Updated Sections 4 and 6
0.6	01/12/23	Dimitris Ntalaperas	UBI	Updated Section 3
0.7	03/12/23	Andreas Karabetian	UPRC	Preparation for internal review
0.8	06/12/23	Georgia Panagiotidou	AUTH	Reviewed document
0.9	12/12/23	Vicky Stroumpou	MT	Reviewed document
1.0	28/12/2023	Spiros Borotis	MAG	Final editing and submission

Table of Contents

Abstract	6
1 Introduction	7
1.1 Purpose and scope.....	7
1.2 Document structure.....	7
1.3 Updates with respect to previous version	8
2 Visualization Workbench	9
2.1 Visualization tools for Policy Making and Bias Detection	9
2.2 Architecture	9
2.3 Baseline Technologies.....	11
2.4 Source Code and Availability.....	11
2.5 Installation Guide.....	13
3 Policy Recommendation Dashboard	15
3.1 Overview	15
3.2 Policy Recommendation Toolkit Interface	15
3.2.1 <i>Overview</i>	15
3.2.2 <i>Objectives</i>	18
3.2.3 <i>User Interaction</i>	19
4 Bias Detector Catalogue Dashboard	20
4.1 Overview	20
4.2 Bias Detector Catalogue Interface	20
4.2.1 <i>Overview</i>	20
4.2.2 <i>Objectives</i>	23
4.2.3 <i>User Interaction</i>	23
5 Water Management Dashboards	25
5.1 Overview	25
5.2 Use Cases	25
5.2.1 <i>Drinking Water Forecasting</i>	25
5.2.2 <i>Sewage Water Forecasting</i>	26
6 Sustainable Development Dashboards	28
6.1 Overview	28
6.2 Use Cases	28
6.2.1 <i>Top 100 projects checklist for applicants and reviewers</i>	28
6.2.2 <i>SDG Observatories Methodology Interface</i>	28
6.2.3 <i>Citizen Engagement Using NLP Interface</i>	29
7 Tourism-driven Policy Management and Optimisation Dashboards	30
7.1 Overview	30
7.2 Citizens Feedback Interactive Interface	30
7.3 Use Cases	31
7.3.1 <i>Traffic Violations Interface</i>	31
7.3.2 <i>Waste Management Interface</i>	32
8 Conclusions	34
9 References	35

List of figures

Figure 1: Visualization Workbench Architecture.....	10
Figure 2: Visualization Workbench source code	12
Figure 3: GitLab issues for the Visualization Workbench	13
Figure 4: Homepage of the Policy Recommendation Toolkit	16
Figure 5: Policy Selector interface	16
Figure 6: KPI / constraints interface	17
Figure 7: Recommendation results	17
Figure 8: Compare policies interface.....	18
Figure 9: Comparative analytics	18
Figure 10: Use case diagrams for user interaction with the Policy recommendation toolkit.....	19
Figure 11: Scrollytelling section example 1	21
Figure 12: Scrollytelling section example 2	21
Figure 13: Stages of training section	22
Figure 14: Real examples of bias section.....	22
Figure 15: Real example of bias opened.....	23
Figure 16: Use case diagram for user interaction with the Bias Detector Catalogue	24
Figure 17: Drinking water forecasting interface	26
Figure 18: Sewage water forecasting interface	27
Figure 19: Policies Chatbot interface.....	29
Figure 20: Citizens Feedback interface.....	31
Figure 21: Traffic violations interface.....	32
Figure 22: Waste management interface	33

List of Tables

Table 1: API Endpoint of Visualization Workbench Server	14
---	----

Abbreviations

Abbreviation	Description
AI	Artificial Intelligence
API	Application Programming Interface
BOD5	5-day Biochemical Oxygen Demand
COD	Chemical Oxygen Demand
DPB	Diputación Provincial de Badajoz
IRCAI	International Research Centre On Artificial Intelligence
JSI	Jozef Stefan Institute
NT	Nitrogen Compounds
OECD	Organisation for Economic Cooperation and Development
SDG	Sustainable Development Goal
TP	Total Phosphorus
UNESCO	United Nations Educational, Scientific and Cultural Organization
VVV	Vari-Voula-Vouliagmeni

Abstract

The Visualization Workbench, the central frontend component of the AI4Gov project, is a transformative web application that transcends traditional boundaries, ushering in a new era of data-driven governance. The Workbench provides a robust, scalable, and responsive platform. Hosting seven distinct use cases across three visionary partners in the sectors of Sustainable Development, Tourism, and Water Management, it serves as a dynamic nexus where diverse expertise converges to generate actionable insights. With an emphasis on technology and ethical AI deployment, the Workbench integrates exploitable services like the Bias Detector Toolkit and the Policy Recommendation Toolkit. This tool explores the pivotal role of visualization services in the realms of Artificial Intelligence (AI) policy formulation and bias detection, highlighting their crucial importance in steering ethical and responsible AI practices. In the context of policy formulation, visualization services serve as indispensable tools, facilitating a comprehensive understanding of complex data sets and fostering informed decision-making among policymakers. The ability to visualize intricate relationships and patterns empowers stakeholders to navigate the nuanced landscape of AI governance, promoting transparency and accountability. Concurrently, in the realm of bias detection, visualization services play a critical role in surfacing latent biases within AI algorithms. By rendering biases visually accessible, these services enable a deeper comprehension of algorithmic decision-making processes, fostering a proactive approach to identify, understand, and mitigate biases. This abstract underscores the symbiotic relationship between visualization services and ethical AI practices, emphasizing their collective potential to shape policies that align with fairness, equity, and societal well-being.

1 Introduction

1.1 Purpose and scope

The purpose of this Deliverable is to present the Visualization Workbench component, a dynamic web application that 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, the Visualization Workbench stands as a testament to the convergence of cutting-edge technology and real-world challenges. At its core, this web application functions as the hosting ground for all pilot partner use cases, weaving together the collective wisdom of three visionary partners. These partners, strategically active in the sectors of Water Management, Sustainable Development and Tourism, bring a wealth of domain-specific knowledge to the table.

The Visualization Workbench proudly accommodates seven distinct use cases, each a unique exploration into the potential synergies of AI and governance. Beyond its role as a showcase for these pilot projects, the Workbench extends its influence through two exploitable services: the Bias Detector Toolkit and the Policy Recommendation Toolkit. These services add an extra layer of functionality, addressing critical aspects of fairness and informed decision-making within the realm of governance.

1.2 Document structure

The deliverable is structured as follows:

- **Chapter 1** introduces the deliverable's purpose and scope, its document structure along with any updates compared to the previous version.
- **Chapter 2** presents 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 a video user guide and installation instructions.
- **Chapter 3** describes the Policy Recommendation Dashboard, the front-end application of the Policy Recommendation Toolkit. In this chapter we take a brief overview of the component and interface, the key objectives as well as a user interaction diagram.
- **Chapter 4** provides information on the Bias Detector Catalogue Dashboard, with a brief overview of the component and interface, the key objectives as well as a user interaction diagram.
- **Chapter 5** focuses on the first sector of use cases, Water Management, while briefly describing the corresponding pilot partner. Then it delves deeper into each use case description accompanied by useful screenshots.
- **Chapter 6** focuses on the second sector of use cases, Sustainable Development, while briefly describing the corresponding pilot partner. Then it delves deeper into each use case description accompanied by useful screenshots.

- **Chapter 7** focuses on the third sector of use cases, Tourism-driven Policy Management and Optimisation, while briefly describing the corresponding pilot partner. Then it delves deeper into each use case description accompanied by useful screenshots.
- **Chapter 8** concludes this deliverable, summarising the key findings and next steps.
- **Chapter 9** includes the reference list.

1.3 Updates with respect to previous version

This is the first of the two versions of the deliverable on Visualization Services. The second version, i.e., D4.4 Policies Visualization Services V2, will be delivered on M27 (March 2025).

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, and a comprehensive state-of-the-art analysis that illuminates its significance within the evolving landscape of data visualization. The integration of advanced visualization tools into governance frameworks has become increasingly important for enhancing decision-making processes. The sustainable development and deployment of AI in the public sector require deliberation between developers, decision makers, deployers, end users, and the public (leikas, 2022). The long-term success of AI in the public sector relies on effectively embedding it into delivery solutions for policy implementation (Mikhaylov, 2018), but it is clear that delivering on the AI grand challenge requires collaboration between universities and the public and private sectors.

The Visualization Workbench aligns seamlessly with this perspective, offering a user-centric interface that empowers decision-makers across various sectors. Furthermore, the role of bias detection tools in decision support systems is crucial, as highlighted by (Ltifi, 2016), who proposed visualization tools for intelligent decision support systems to analyse large volumes of data. Additionally, the work of (Law, 2020) focuses on designing tools for semi-automated detection of machine learning biases, addressing the significant financial and ethical implications of undetected biases, a feature prominently integrated into the Visualization Workbench through the Bias Detector Toolkit. This addition not only ensures the ethical use of AI but also aligns with the broader discourse on responsible and fair AI implementation in governance.

In addition, authors in (Wu, 2020) demonstrate the application of an interactive web-based geovisual analytics platform for co-clustering spatio-temporal data, allowing the concurrent analysis of spatial patterns, something that the Visualization Workbench also aims to achieve. As we navigate the state-of-the-art landscape, it becomes evident that the Visualization Workbench is not just a standalone component but a dynamic response to the evolving needs of governance, drawing from and contributing to the cutting-edge research in the field.

2.2 Architecture

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

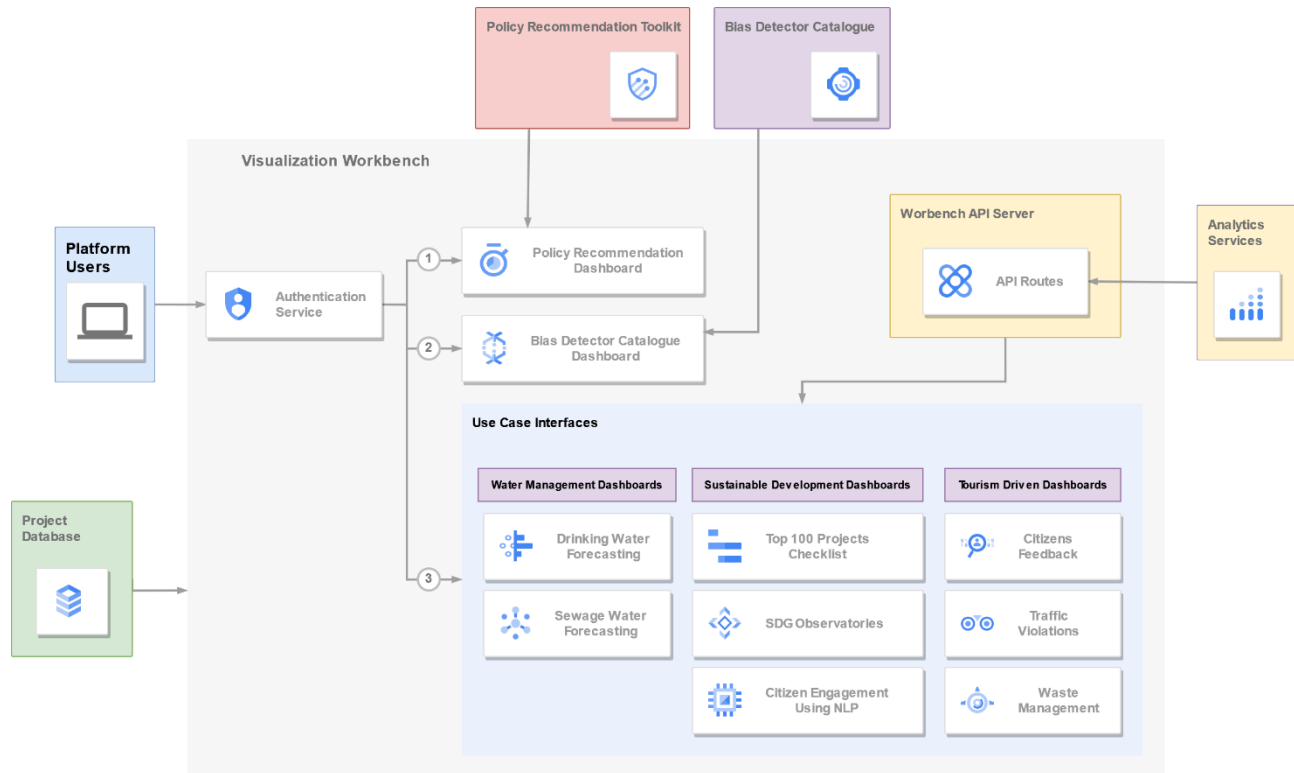


Figure 1: Visualization Workbench Architecture

The architecture of the Visualization Workbench is designed to provide a seamless and interactive experience for users, integrating authentication, specialized interfaces, and communication channels between various components of the AI4Gov project. At its core is the web application, forming the user-facing front end. Users initiate their journey by authenticating themselves, after which they are presented with a multifaceted interface offering three key options: the Bias Detector Catalogue interface, the Policy Recommendation Toolkit interface, or access to one of the seven distinct use case interfaces.

The communication backbone of the Visualization Workbench is facilitated by an API server. This server acts as the intermediary, ensuring smooth interaction between the user interface and the other components that constitute the AI4Gov project. The project's databases play a pivotal role, serving as repositories for metadata and static content. These databases, seamlessly integrated into the architecture, contribute to the persistence and retrieval of critical information, ensuring that user interactions and project data are stored efficiently and securely.

The Bias Detector Catalogue interface and the Policy Recommendation Toolkit interface showcase the specialized services offered by the Visualization Workbench. These components play a crucial role in addressing bias within AI models and providing intelligent policy recommendations, respectively. Each of the seven use case interfaces represents a unique sector-specific exploration, fostering a tailored experience for users engaged in Sustainable Development, Tourism, and Water Management.

Overall, the architecture of the Visualization Workbench harmonizes user interaction, API communication, and database utilization to create a comprehensive platform that not only caters to the diverse needs of the AI4Gov project but also prioritizes a user-friendly and efficient experience.

2.3 Baseline Technologies

The backbone of the Visualization Workbench component is a robust technological stack carefully selected to ensure efficiency, scalability, and seamless integration. At its core, the Workbench is built as a Next.js (Next.js, n.d.) web application, harnessing the power of this React (React, n.d.) 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 a Node.js (Node.js, n.d.) API server, leveraging the asynchronous, event-driven architecture of Node.js 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 both the Next.js application and Node.js server.


This containerization ensures that the Workbench's components, with all their 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.



AI4Gov > T4.4 - Visualization Workbench




T4.4 - Visualization Workbench  Project ID: 1 








🔔  ☆ Star 0 🍴 Fork 0


🔗 31 Commits 🌿 1 Branch 🏷️ 0 Tags 📄 14.9 MB Project Storage

This project contains all the source code that is developed in the context of T4.4 - Interactive Self-Explained Visualization Workbench.


 **laptop updates and dockerfile** adreaskar authored 2 weeks ago ab9d20e1 

main visualization-workbench /  Find file Web IDE  Clone 

 README  Auto DevOps enabled  Add LICENSE  Add CHANGELOG  Add CONTRIBUTING  Add Kubernetes cluster  Add Wiki

 Configure Integrations

Name	Last commit	Last update
📁 application	laptop updates and dockerfile	2 weeks ago
📄 README.md	updated readme	4 months ago

 README.md

T4.4 - Visualization Workbench

Integrate with your tools

Figure 2: Visualization Workbench source code

There have also been created several issues in order to keep track of all use case requirements and maintain an organized repository and project.

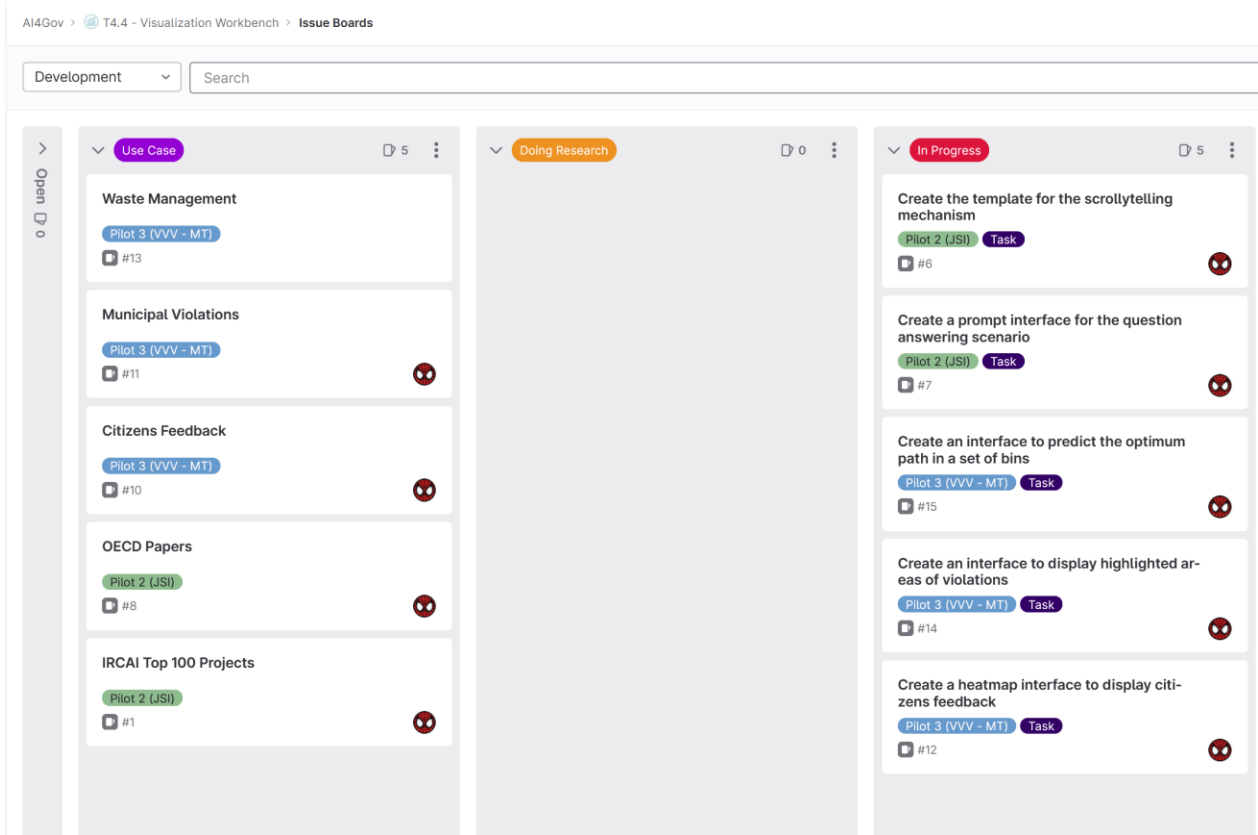


Figure 3: GitLab issues for the Visualization Workbench

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. Once the containers start, the front-end can be found on port 3000 and the backend server on port 8281.

At the moment the backend server only utilizes one API endpoint, as described in Table 1.

Endpoint Name	ai4gov_chatbot_api
Endpoint URL	http://[SERVER IP]:PORT/api/chatbot
Description	This endpoint is responsible for interacting with the analytics services regarding the Policies Chatbot, in order to send the user input and retrieve the analysed answer.
HTTP Method	POST
Parameters	prompt: The user prompt

Response	JavaScript object containing two fields. 1. answer: The short answer that the analytics service returns 2. context: The context in which the answer was found
-----------------	---

Table 1: API Endpoint of Visualization Workbench Server

A video demonstrating all the functionalities and interfaces can be found [here](#). The implementation of the Visualization Workbench is on-going, which means that the functionalities presented in this deliverable are subject to change, as the project progresses and other user requirements might surface from the pilot partners. In any case, in the second version of this deliverable, i.e., D4.4 Policies Visualization Services V2, the finalized version of the implementation should be provided.

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 in 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 still in development and more concrete results can be explored in the D3.3 – Policy Recommendation Toolkit V1, which is due on M16.

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. If more than one domains are available to the policy maker; this can be chosen via the initial page (Figure 4). Then a policy can be chosen (Figure 5) and viewed and/or edited via addition or re-definition of the relevant KPIs (Figure 6). Selecting policies involves the definition of the relevant KPIs (Figure 6). Selecting policies involves the definition of hard and soft constraints imposed for the desired policy; the framework will then retrieve the policies that the end-user can select a set of policies (Figure 8) and view comparative analytics (Figure 9) by using easy to understand visualisations (the form of the visualisations and the level of customisation will depend on the details of the use cases and are still in the design phase).

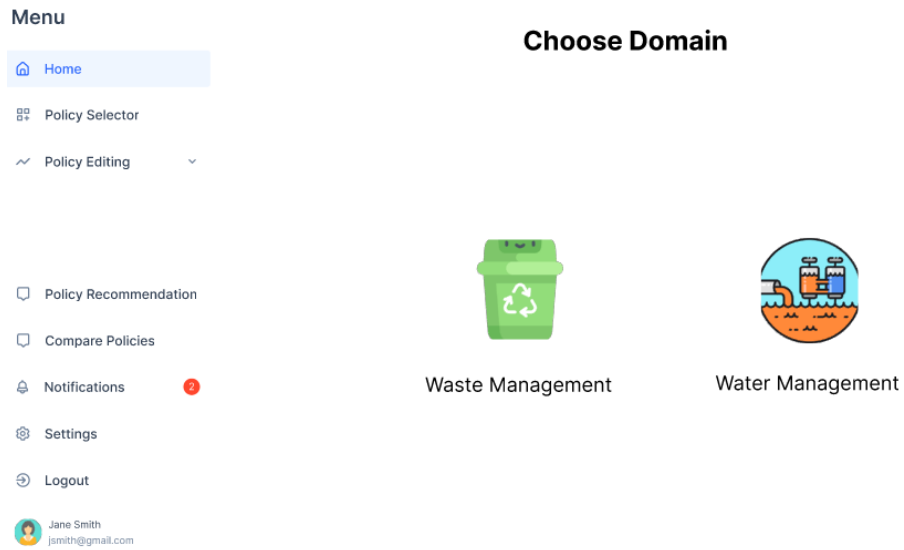


Figure 4: Homepage of the Policy Recommendation Toolkit

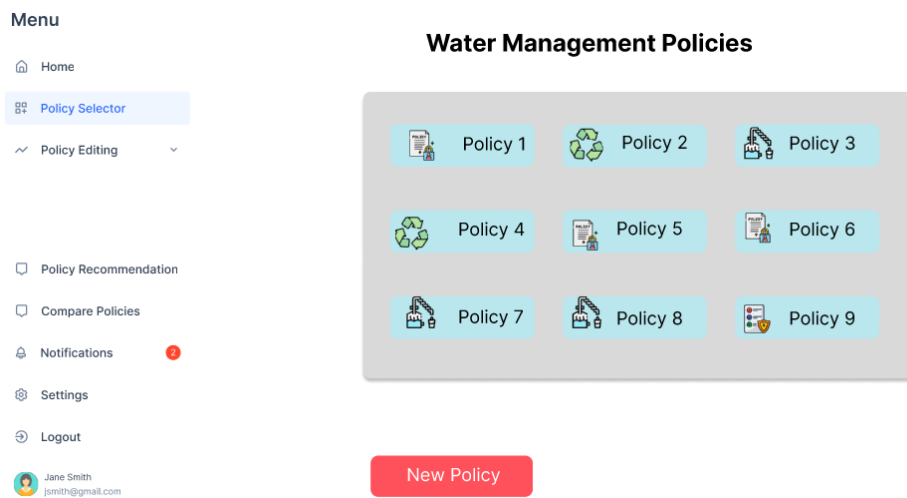


Figure 5: Policy Selector interface

Menu

- Home
- Policy Selector
- Policy Editing
- Policy Recommendation**
- Compare Policies
- Notifications
- Settings
- Logout

Jane Smith
jsmith@gmail.com

Set Criteria

Hard Constraints

KPI	Target	Condition	
Cost Reduction	20%	Greater than	-
			+
Water Reclamation			
Water Purity			

Soft Constraints

KPI	Target	Condition	
Approval rate	62%	Greater than	-
			+

Recommend Policy

Figure 6: KPI / constraints interface

Menu

- Home
- Policy Selector
- Policy Editing
- Policy Recommendation**
- Compare Policies
- Notifications
- Settings
- Logout

Jane Smith
jsmith@gmail.com

Recommendation Results

Retrieved Policies

	Hard Constraint 1 Score	Hard Constraint 2 Score	Soft Constraint 2 Score
Policy 1	63%	58%	2%
Policy 2	51%	38%	18%

Legend

- Good
- Pass
- Fail

Back

Figure 7: Recommendation results

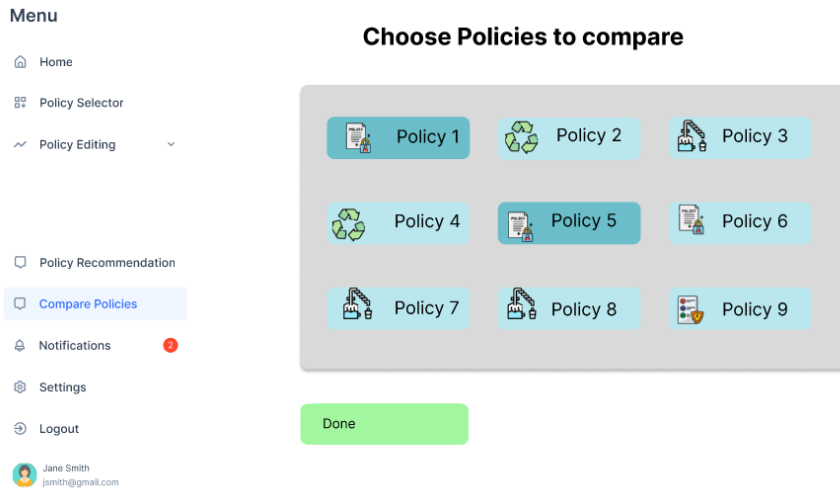


Figure 8: Compare policies interface

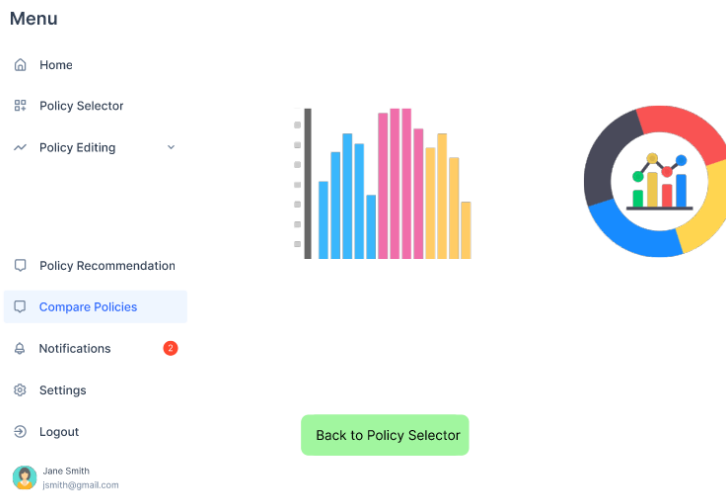


Figure 9: Comparative analytics

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.
- Provide an engine to combine policies and evaluate their output on specific scenarios.
- Interface with the AI4Gov blockchain to execute smart contracts that execute commonly defined and endorsed business rule to evaluate existing or custom policies.

3.2.3 User Interaction

Figure 10 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 backend's smart contract infrastructure.

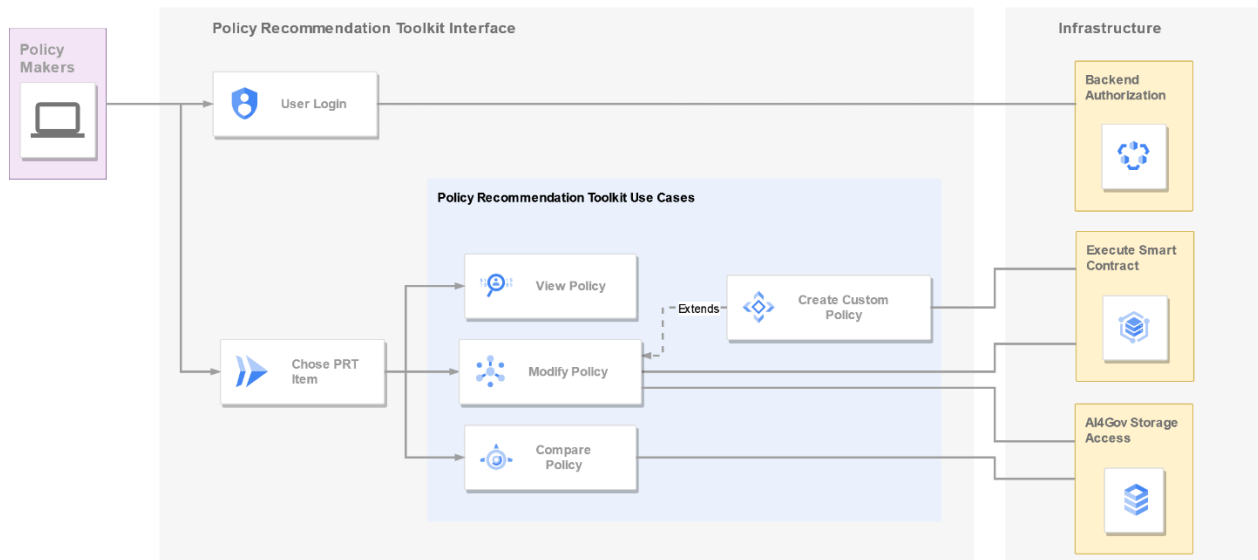


Figure 10: Use case diagrams for user interaction with the Policy recommendation toolkit

4 Bias Detector Catalogue Dashboard

4.1 Overview

As governments increasingly adopt AI services, the repercussions of model decisions reverberate through a growing expanse of people's lives. Hence, it becomes paramount to ensure that models employed in policy contexts do not perpetuate biases inherent in real-life situations and data. In pursuit of this objective, we are crafting the Bias Detector Toolkit. This toolkit is meticulously designed to empower policymakers, researchers, and practitioners with a nuanced understanding of a diverse array of tools tailored for detecting biases in AI systems.


Through the incorporation of visual synthesis, the toolkit not only augments accessibility but also cultivates a more informed and effective approach to bias detection. This strategic approach harmonizes seamlessly with our commitment to bridging the realms of data science and policy formulation, providing valuable insights to navigate the complex intersection of these domains.

4.2 Bias Detector Catalogue Interface

4.2.1 Overview

Bias Detector Toolkit that serves as a visual catalogue synthesizing various bias mitigation tools for bias detection in AI. The toolkit is designed to empower general public, developers, researchers, and practitioners with a nuanced understanding of the diverse range of tools available for detecting and mitigating biases in AI systems. The toolkit is unveiled through scrollytelling (Figure 11 and Figure 12), a narrative method intertwining scrolling and storytelling. A section where the stages of training a machine learning model is also explored (Figure 13), while also mentioning the potential bias that can occur on each step. Using vivid visual metaphors, we navigate the intricacies of bias, emphasizing the need for robust mitigation strategies.

Providing real life examples (Figure 14 and Figure 15), all relevant stakeholders can get a better grip of importance for bias mitigation. Structured as a dynamic visual synthesis, the catalogue offers a comprehensive overview of bias mitigation tools, categorized by functionalities and applications. Serving as a visual catalogue, it facilitates exploration, comparison, and informed tool selection, fostering a more effective approach to bias mitigation.



Bias Detector Catalogue



The aim of this Use Case is to create a document-style framework with listed services and a set of rules for bias evaluation. In the long run, the UC will establish a clearinghouse for bias evaluation and have a model, rules, and toolkit for bias evaluation of any AI solution in general.



In our everyday lives, we often encounter bias—a concept deeply ingrained in human society. Bias refers to the presence of unfair and prejudiced attitudes or behaviors that favor one group over another, whether it's based on race, gender, age, or socio-economic status. These biases can have profound consequences, like gender pay gaps or racial disparities in access to opportunities. They're not always intentional; they can also result from unconscious or systemic factors perpetuating inequality.

Now, consider the remarkable world of Artificial Intelligence (AI). AI represents a technological leap, enabling machines to perform tasks requiring human intelligence. Think of voice assistants like Siri or Alexa, which recognize and respond to your voice commands, or complex systems aiding medical diagnoses or autonomous vehicles. AI has the potential to revolutionize industries and improve our lives, but it also raises

Figure 11: Scrollytelling section example 1



Consider a scenario where a popular e-commerce platform employs AI-driven product recommendation algorithms. These algorithms aim to enhance the user experience by suggesting products based on past purchases and browsing history. However, if the training data used for these algorithms predominantly reflects biased patterns, it can lead to problematic recommendations. For instance, if the historical data shows a preference for gender-specific products due to societal stereotypes, the AI might inadvertently suggest products reinforcing those stereotypes. This means that a woman searching for tech gadgets might receive recommendations for makeup items simply because the AI infers her gender from the data. In such cases, bias in AI not only limits product diversity but also perpetuates gender stereotypes, creating a less inclusive shopping experience and reinforcing societal biases.

Consider the case of an AI-powered criminal sentencing tool used by some judicial systems to predict the likelihood of reoffending. This system relies on historical data about past convictions, sentences, and reoffending rates to make recommendations for sentencing decisions. The problem arises when this data reflects existing biases in the criminal justice system, such as racial disparities in arrests and sentencing. If the training data contains these biases, the AI model may learn to perpetuate them. Consequently, the tool might unfairly recommend longer sentences for individuals from certain racial or socio-economic backgrounds, exacerbating the very disparities it was meant to address. This real-world example vividly illustrates how bias in AI can have profound and unjust consequences in areas as critical as criminal justice, highlighting the urgency of

Figure 12: Scrollytelling section example 2

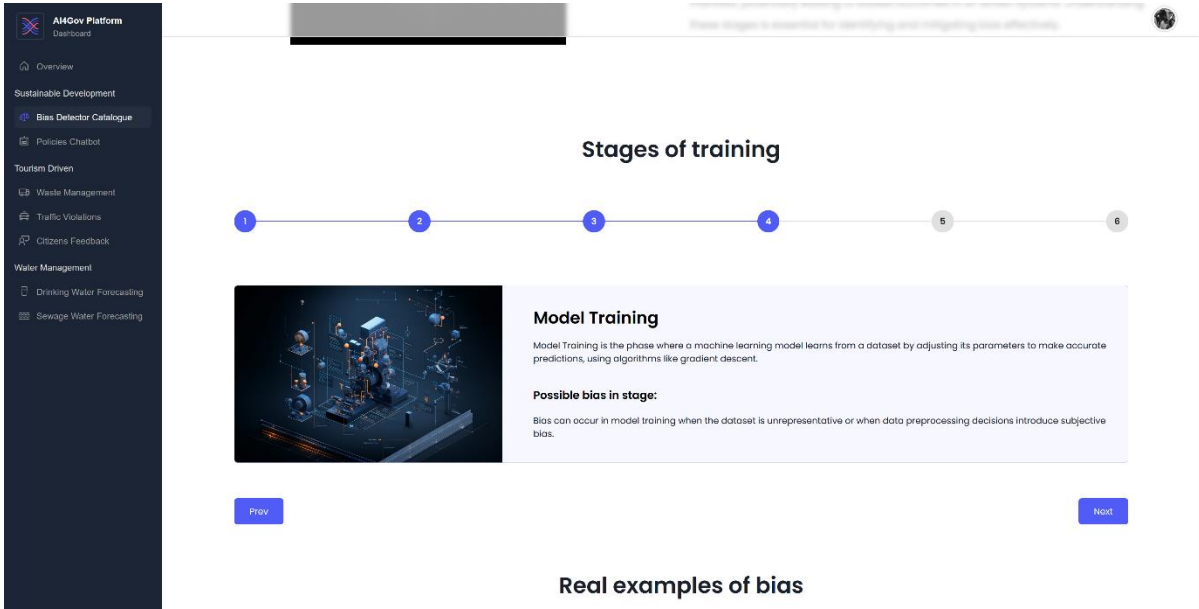


Figure 13: Stages of training section

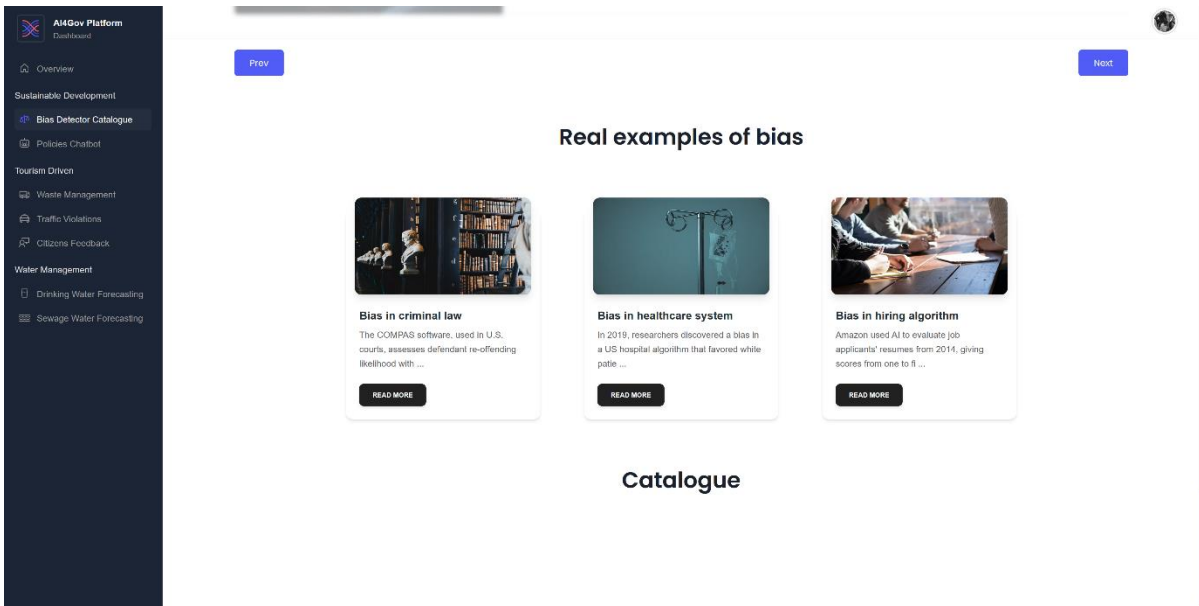


Figure 14: Real examples of bias section

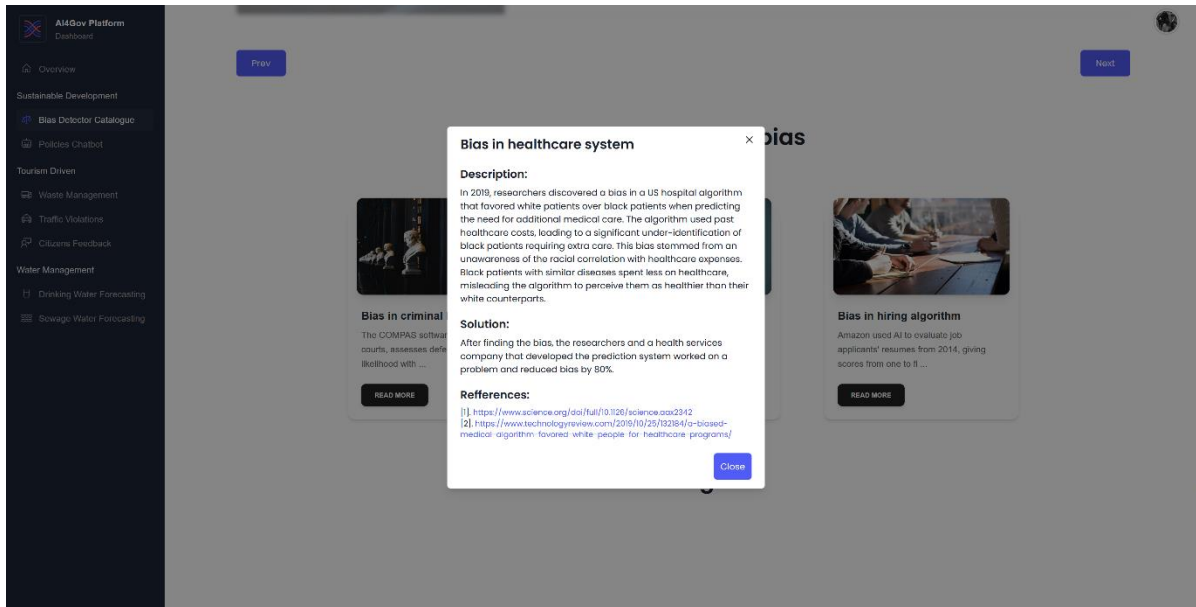


Figure 15: Real example of bias opened

4.2.2 Objectives

This tool aims to achieve 3 different goals, aiming to 3 different audiences.

- For the general public and policymakers, our innovative strategy employs scrollytelling and real-life examples, paving the road to a nuanced understanding of the critical importance of bias mitigation.
- For AI/ML engineers immersed in the technical intricacies, our road leads to a meticulously curated catalogue. This catalogue, enriched by visual synthesis, expedites access to a comprehensive array of currently available tools and methods for bias detection and mitigation. By structuring the catalogue based on functionalities and applications, we facilitate a streamlined approach for engineers, enabling them to navigate the complex landscape of bias mitigation tools with efficiency and precision.
- For specific users from use cases (e.g. Top100 projects applicants and reviewers; SDG observatories developers, potentially also developers in other two pilots), additional layer will be revealed by login into platform, where they will be able to access additional layer. E.g., for applicants and reviewers to Top100 projects, checklist will be provided.

4.2.3 User Interaction

User interaction with our platform is designed to be intuitive, seamless, and tailored to meet the diverse needs of stakeholders. Through a user-friendly interface, individuals can effortlessly navigate the platform, accessing relevant tools and resources based on their roles and objectives.

Upon accessing the platform, users are greeted with a clear and organized layout, facilitating easy exploration and interaction. If they follow along, they are presented with the scrollytelling

explanation of bias and bias mitigation in AI, followed by real life examples, and strategies employed previously. Lastly, Bias Detector Catalogue is presented, allowing dynamic exploration. Alternatively, the user can skip the scrollytelling and real-life examples explanation part and jump directly to The Bias Detector Catalogue (we expect this with often use of the tool).

For specific user groups, such as Top 100 project applicants and reviewers or developers associated with SDG observatories, the platform offers tailored layers of functionality. Secure login credentials unlock additional features, providing a customized experience based on the user's specific needs and roles. Whether it's accessing checklists, specialized tools, or collaborative spaces, the platform ensures that users can engage with relevant resources efficiently and effectively.

Overall, our platform prioritizes user-centric design, fostering a positive and productive interaction that aligns with the unique requirements of stakeholders.

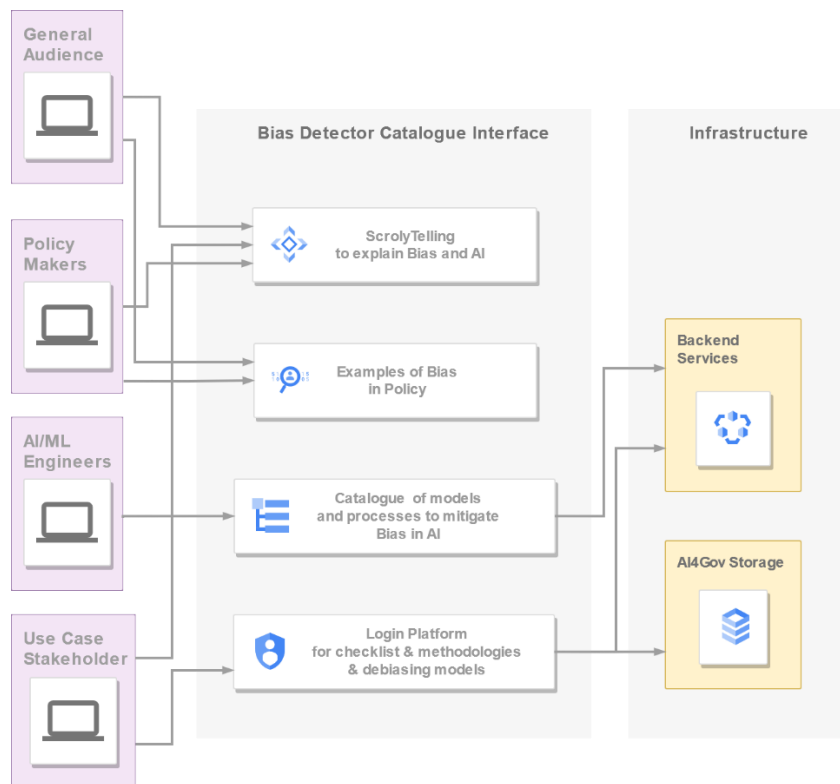


Figure 16: Use case diagram for user interaction with the Bias Detector Catalogue

5 Water Management Dashboards

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

This user interface is designed to facilitate the seamless input and analysis of drinking water data through an intuitive form. The form includes two crucial components: a time range specification and a variable input for the drinking water parameter, such as pH levels. Users are prompted to define a date range by specifying the start date (date from) and end date (date to), ensuring a comprehensive time frame for data analysis. Additionally, users input the specific variable of interest related to drinking water quality, exemplified by pH, which is fundamental for understanding the water's acidity or alkalinity.

Once the user submits this information, the analytics services harnesses advanced time series forecasting algorithms to generate a predictive chart. This chart provides a visual representation of the anticipated future values of the selected variable within the defined time range. The forecasting model takes into account historical data within the specified timeframe, identifying patterns, trends, and potential fluctuations. The resulting chart empowers users with valuable insights into the expected behaviour of the chosen drinking water parameter, aiding in proactive decision-making for water quality management and resource planning. The user-friendly nature of the interface, coupled with the sophisticated forecasting capabilities, positions it as a powerful tool for making informed decisions regarding drinking water quality based on historical trends and predictive analytics.

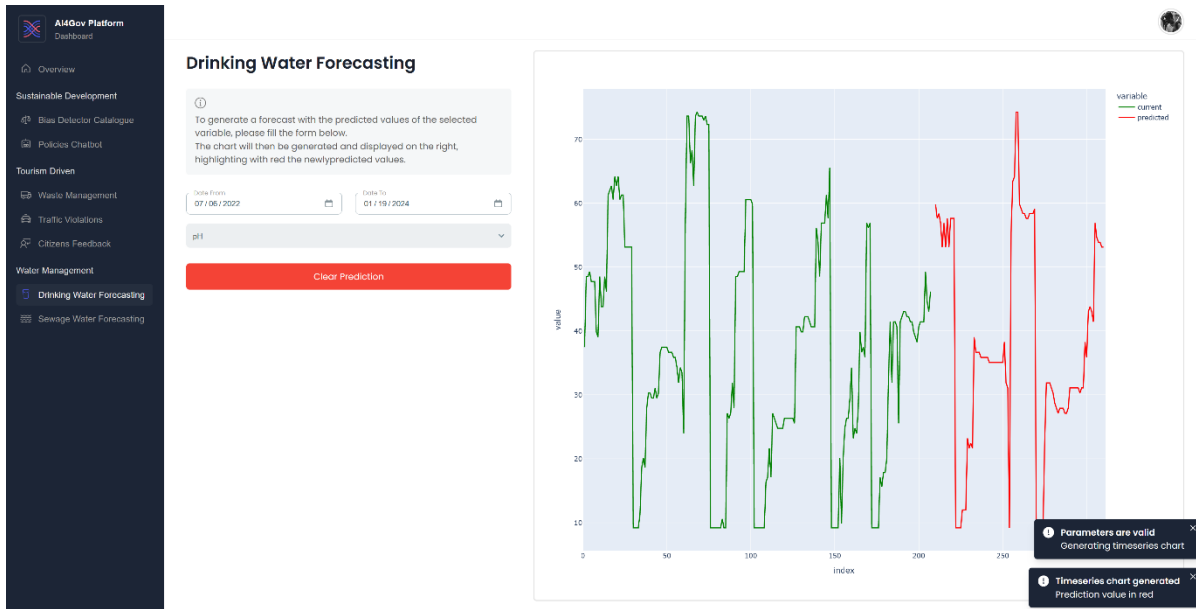


Figure 17: Drinking water forecasting interface

5.2.2 Sewage Water Forecasting

This user interface offers a robust platform for the input and analysis of sewage water variables, with a focus on parameters critical to wastewater quality assessment, such as chemical oxygen demand (COD), total phosphorus (TP), nitrogen compounds (NT), and 5-day biochemical oxygen demand (BOD5). The form within the interface seamlessly guides users through the process, presenting two key components: a flexible time range specification and variable inputs for sewage water parameters. Users are prompted to define a date range by indicating the start date (date from) and end date (date to), ensuring a tailored temporal scope for thorough data exploration. Simultaneously, users input specific sewage water variables of interest, encompassing parameters like COD, TP, NT, and BOD5, pivotal for assessing the organic and inorganic components of wastewater.

Upon submission, the analytics services employ advanced time series forecasting algorithms to generate an insightful predictive chart. This visual representation offers a glimpse into the anticipated future values of the selected sewage water variables within the defined time range. The forecasting model harnesses historical data from the specified timeframe, deciphering complex patterns, trends, and potential fluctuations in the concentrations of COD, TP, NT, and BOD5. The resulting chart empowers users with actionable foresight, facilitating proactive decision-making in wastewater treatment and environmental management. The user-friendly interface, coupled with cutting-edge forecasting capabilities, positions this tool as an invaluable resource for making informed decisions about sewage water quality, leveraging historical trends and predictive analytics across a spectrum of essential parameters.

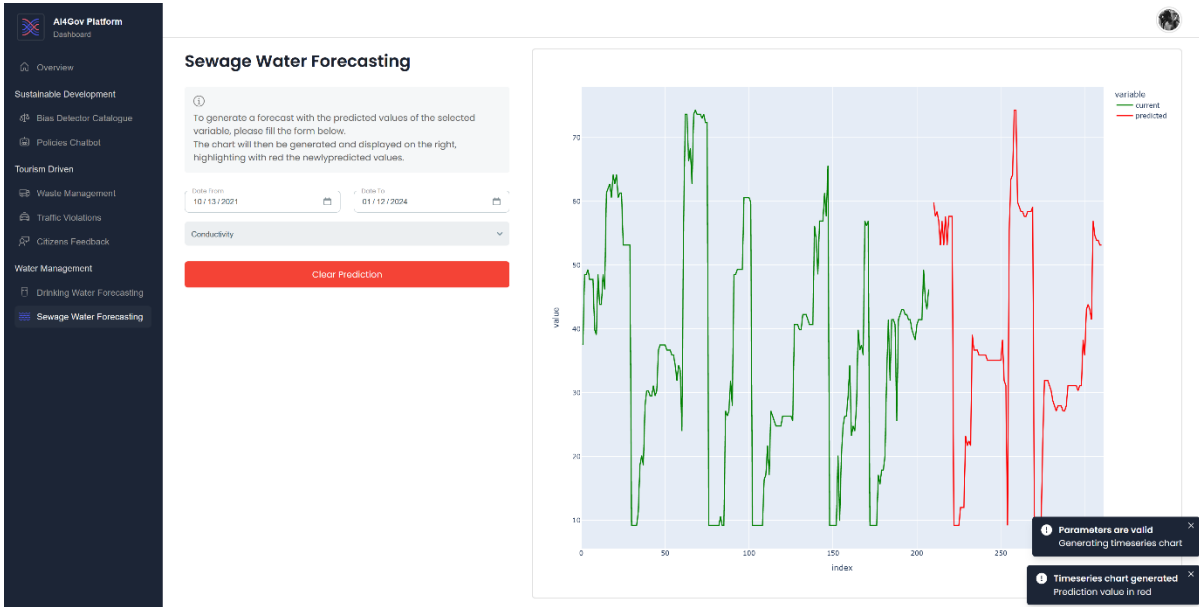


Figure 18: Sewage water forecasting interface

6 Sustainable Development Dashboards

6.1 Overview

International Research Centre in Artificial Intelligence (IRCAI), under the wing of JSI, is an organisation 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 will be checklists both for applicants and reviewers. Applicants to the Top100 projects will be 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, will be similarly presented with a checklist upon review of application, that will enable them to evaluate to what extent bias component has been considered in the models of the Top100 applications.

As this use case is in early stages, more details will be presented in the next version of this Deliverable.

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 are under development. Contribution of visualization workbench to this use case is to develop visual markings, that will enable tagging some datasets, that are more prone to bias, and that have potentially missing data.

As this use case is in early stages, more details will be presented in the next version of this Deliverable.

6.2.3 Citizen Engagement Using NLP Interface

Introducing the Policies Chatbot Interface, a platform designed to facilitate comprehensive insights into AI advancements across various countries. This innovative interface operates seamlessly as a chatbot, allowing users to effortlessly select a specific country of interest and pose inquiries through a user-friendly prompt. Users can delve into the vast repository of OECD AI papers related to the chosen country, extracting valuable insights and information. Upon receiving a user's question, the chatbot undertakes a sophisticated analysis of the relevant papers, employing natural language processing and advanced data mining techniques to distil concise and accurate responses. The user experience is enhanced by the chatbot's ability to not only provide straightforward answers but also to furnish contextual details gleaned from the wealth of information within the papers.

This dual functionality ensures that users not only receive immediate and comprehensible responses but also gain a deeper understanding of the broader landscape of AI research and development within the selected country. Whether it's inquiries about AI policy frameworks, technological advancements, or societal impacts, the Policies Chatbot stands as a powerful tool, bridging the gap between complex research and user-friendly accessibility.

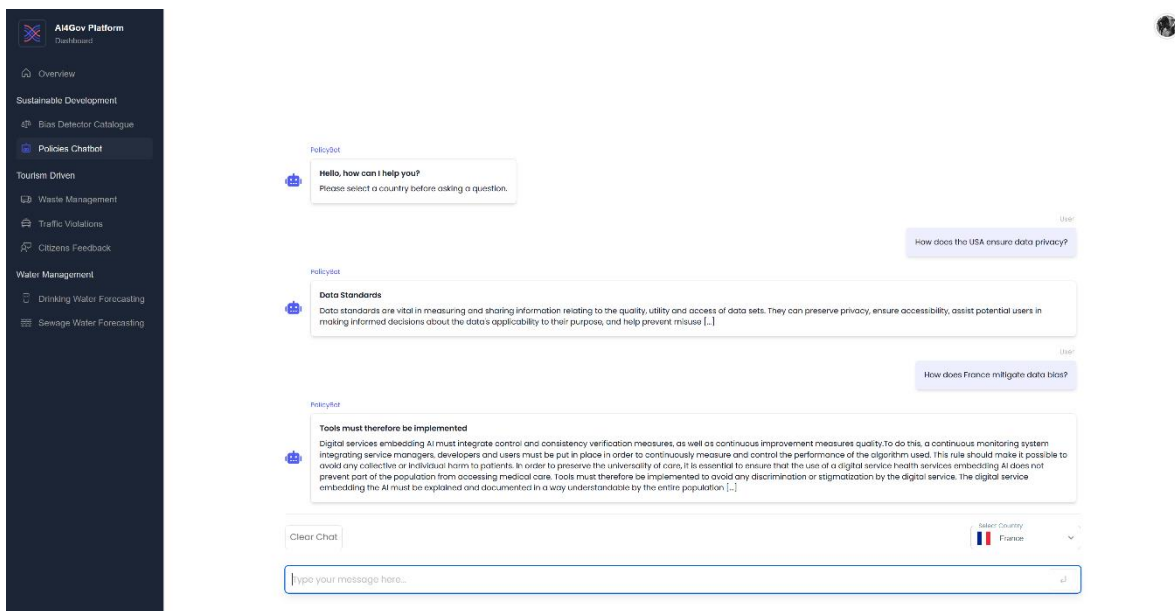


Figure 19: Policies Chatbot interface

7 Tourism-driven Policy Management and Optimisation Dashboards

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 flow of citizens, 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 of citizens 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 Citizens Feedback Interactive Interface

This user interface is designed to provide a seamless experience for tracking various issues within the municipality of Vari - Voula - Vouliagmeni. The focal point of this interface is a comprehensive form where users can input specific details related to the incident they wish to track. The form prompts users to define a precise time range by selecting both a 'date from' and a 'date to,' offering granularity in pinpointing when the issue occurred or persisted. Additionally, users are required to classify the nature of the issue from a predefined list that encompasses a spectrum of concerns such as potholes, environmental damage, overgrown vegetation, damaged pavements, and more.

Once the user submits this information, the analytics services harness the power of dynamic data visualization to generate an interactive heat map. This heat map dynamically illustrates the spatial distribution and intensity of recorded incidents pertaining to the specified issue type within the designated time frame. The color gradient on the heat map serves as an intuitive visual aid, with darker hues representing higher concentrations of reported issues. Users can easily interpret the geographical hotspots of specific problems, enabling both the public and municipal authorities to gain valuable insights into patterns, trends, and priority areas for intervention. This user-friendly and informative interface not only facilitates efficient issue reporting but also empowers local communities and governing bodies with actionable data to enhance the overall management and upkeep of the municipality's premises.

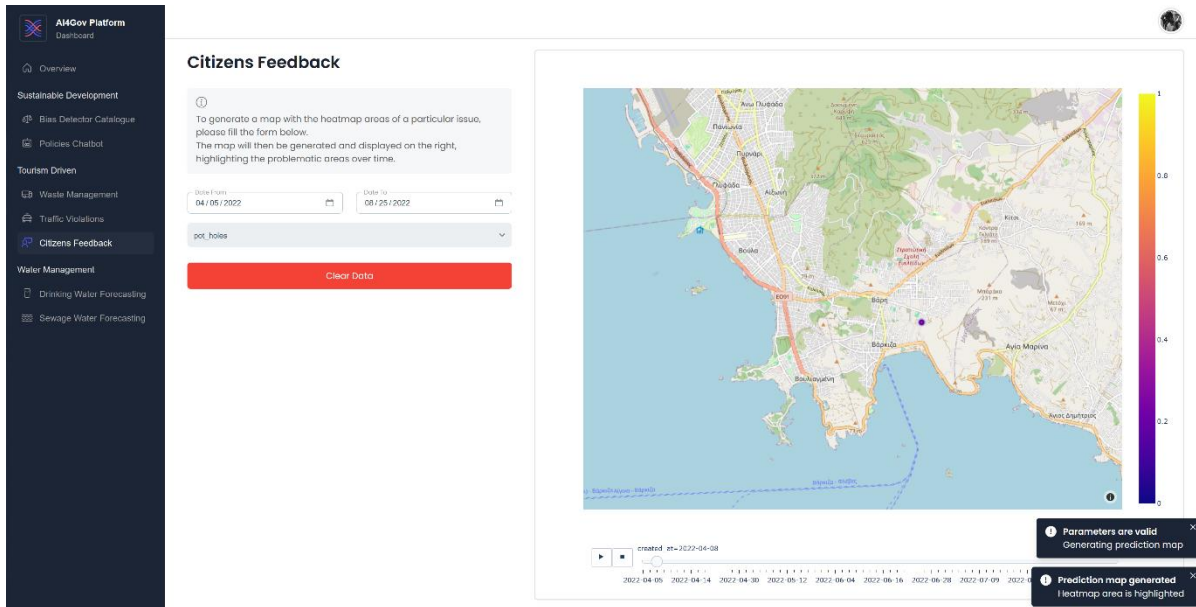


Figure 20: Citizens Feedback interface

7.3 Use Cases

7.3.1 Traffic Violations Interface

The Municipality faces a significant problem with Parking especially during the summer and weekends due mainly to the flow of visitors and tourists. The Municipal administration seeks to handle the problem of Parking and needs a tool that monitors the number of parking 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 premises of 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.

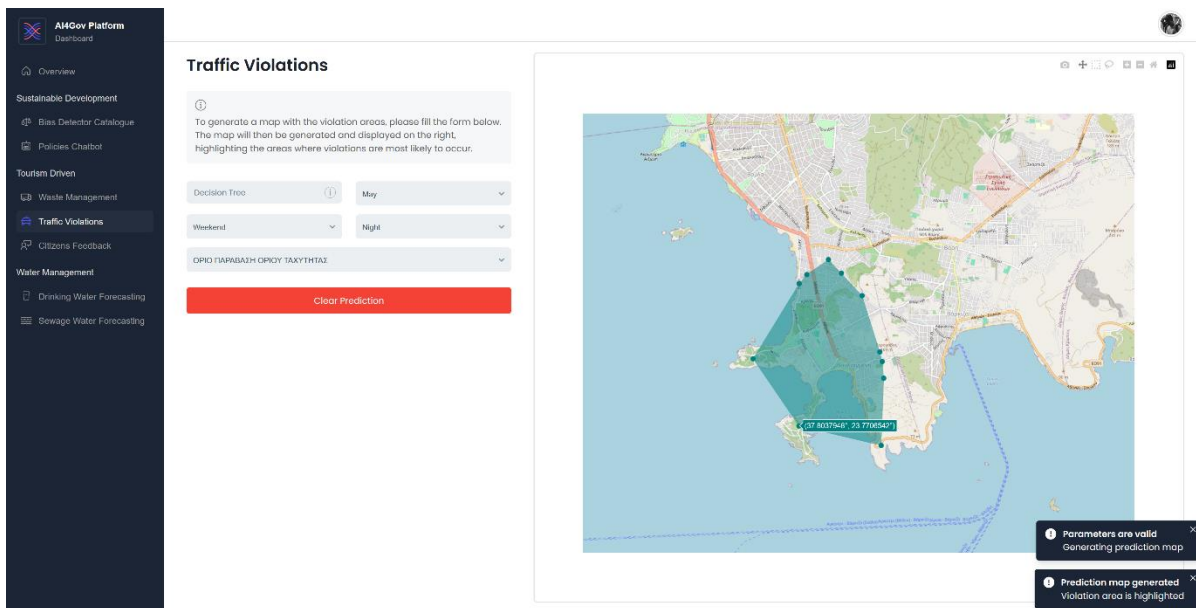


Figure 21: Traffic violations interface

7.3.2 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 staff and resources. Furthermore, as it seeks to implement the *Pay As You Throw* System to whole area of the Municipality, it needs a predictive tool as to the outcome of this expansion in economic terms.

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 be Brown for organic waste, Blue for recyclables, 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 liters) 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 premises. 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.

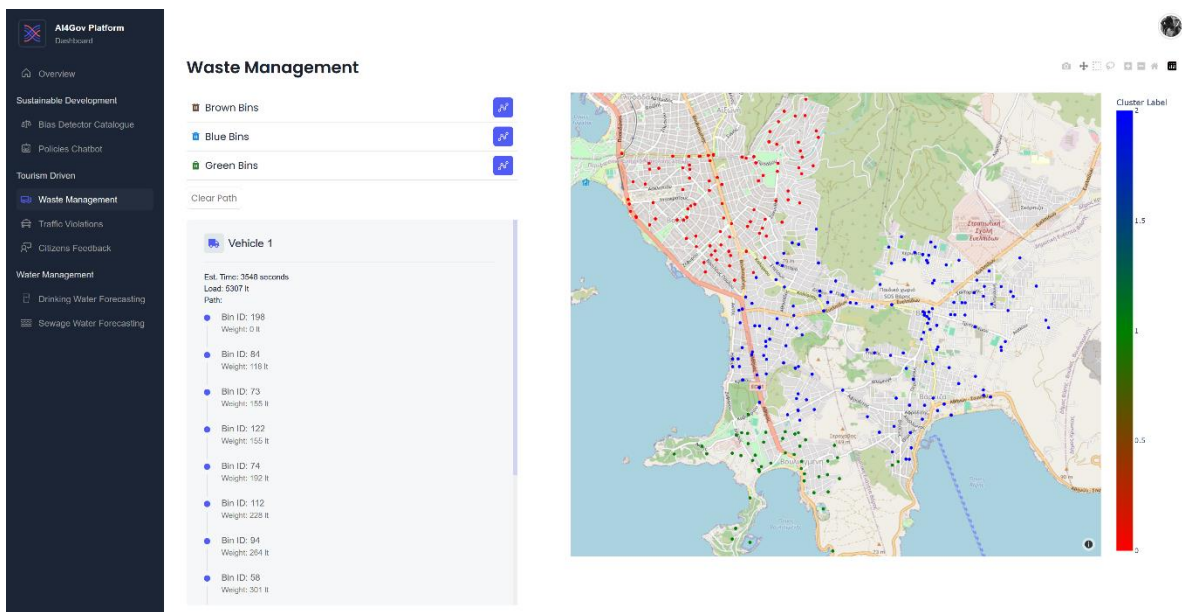


Figure 22: Waste management interface

8 Conclusions

In conclusion, the Visualization Workbench embodies the transformative potential of AI within the context of diverse sectors. Its role as the front end of the AI4Gov project is pivotal, serving as a gateway to seven distinct use cases spanning Sustainable Development, Tourism, and Water Management, each a testament to the collaborative efforts of forward-thinking partners. The Workbench not only hosts these use cases but enriches their impact through two exploitable services — the Bias Detector Toolkit and the Policy Recommendation Toolkit — emblematic of its commitment to ethical AI and informed decision-making.

As the nexus of data visualization and governance, the Visualization Workbench not only streamlines complex datasets but also fosters a holistic understanding of multifaceted challenges. Its journey is not just a testament to the present capabilities of AI but a promise of continuous evolution, adaptability, and responsiveness to the ever-changing landscape of governance challenges and opportunities.

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/>.
- Node.js*. (n.d.). Retrieved from <https://nodejs.org/en>.
- React*. (n.d.). Retrieved from <https://react.dev/>.
- 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>