



Two distributed models implementing the ARSINOE Dynamic Multi-System Resilience and Modelling Framework

Deliverable 3.9

WP3: Dynamic Multi-Sectoral Resilience Modelling and Assessment Framework

Authors: Conor Crilly (UEXE), Navonil Mustafee (UEXE), Dimitris Kofinas (UTH), Otto Chen (UEXE), Barry Evans (UEXE), Lydia Vamvakeridou-Lyroudia (KWR&UEXE), Sophia Papageorgiou (ADDMA), Giannis Adamos (UTH), Giorgos Papangelis (NOA), Albert Chen (UEXE)

Date: 31/03/25



This project has received funding from the European Union's Horizon H2020 innovation action programme under grant agreement 101037424.

Deliverable Number and Name	D3.9 – Two distributed models implementing the ARSINOE Multi-System Resilience and Modelling Framework
Work Package	WP3 – Dynamic Multi-Sectoral Resilience Modelling and Assessment Framework
Dissemination Level	Public
Author(s)	Conor Crilly (UEXE), Navonil Mustafee (UEXE), Dimitris Kofinas (UTH), Otto Chen (UEXE), Barry Evans (UEXE), Lydia Vamvakeridou-Lyroudia (KWR&UEXE), Sophia Papageorgiou (ADDMA), Giannis Adamos (UTH), Giorgos Papangelis (NOA), Albert Chen (UEXE)
Primary Contact and Email	Conor Crilly, c.crilly@exeter.ac.uk
Date Due	31/03/25 (M42)
Date Submitted	31/03/25
File Name	ARSINOE-D3.9-Two-Distributed-Models
Status	Final Version
Reviewed by (if applicable)	Martin Drews (DTU)
Suggested citation	Crilly, C., Mustafee, N., Kofinas, D., Chen, O., Evans, B., Vamvakeridou-Lyroudia, L.S., Papageorgiou, S., Adamos, G., Papangelis, G. and Chen, A. Two distributed models implementing the ARSINOE Multi-System Resilience and Modelling Framework. ARSINOE Deliverable 3.9, H2020 grant no 101037424.

© ARSINOE Consortium, 2025

This deliverable contains original unpublished work except when indicated otherwise. Acknowledgement of previously published material and of the work of others has been made through appropriate citation, quotation, or both. Reproduction is authorised if the source is acknowledged.

This document has been prepared in the framework of the European project ARSINOE. This project has received funding from the European Union's Horizon 2020 innovation action programme under grant agreement no. 101037424.

The sole responsibility for the content of this publication lies with the authors. It does not necessarily represent the opinion of the European Union. Neither the EASME nor the European Commission are responsible for any use that may be made of the information contained therein.

TABLE OF CONTENTS

LIST OF ABBREVIATIONS	5
LIST OF FIGURES	6
LIST OF TABLES.....	7
EXECUTIVE SUMMARY.....	8
1.0 INTRODUCTION	9
2.0 DISTRIBUTED MODEL 1: A HYBRID SIMULATION MODEL FOR MULTI-SECTOR RESOURCE ALLOCATION DURING FLOOD EVENTS	10
2.1 Case Study Overview.....	10
2.2 Methodology	10
2.2.1 Hybrid Simulation using ABS, DES and SD modelling sub-components	10
2.2.2 Data Sources	11
2.2.3 Agent Definitions and Logic	12
2.3 Results & Experiments	13
2.4 Discussion	16
3.0 DISTRIBUTED MODEL 2: AN AGENT-BASED MODEL FOR RESOURCE ALLOCATION DURING HEATWAVE EVENTS	17
3.1 Case Study Overview.....	17
3.2 Methodology	17
3.2.1 An Overview of Agent-based simulation.....	17
3.2.2 Data Sources (outputs of sub-models used as input for capacity model)	18
3.2.3 Risk.....	21
3.2.4 Heatwave Classification.....	24
3.2.5 Resource Definitions.....	25
3.2.6 Agent Logic.....	25
3.3 Results & Experiments	27
3.4 Discussion	32
4.0 CONCLUSION	33
REFERENCES	34

List of abbreviations

ABS	Agent-based Simulation.
CI	Critical Infrastructure.
CTM	Chemistry Transport Model.
DES	Discrete Event Simulation.
DMRM&AF	Dynamic Multi-System Resilience Modelling & Assessment Framework.
ER	Emergency Response.
ER-DT	Emergency Response-Digital Twin
RAM	Resource Allocation Model.
LSM	Land-Surface Model.
SD	System Dynamics.
SHVI	Socioeconomic Heat Vulnerability Index.
UCM	Urban Canopy Model.
WRF	Weather Research and Forecasting.

List of figures

- Figure 1** Conceptual model of the RAM model.
- Figure 2** Water levels at each critical infrastructure location used in the RAM model.
- Figure 3** Total service capacity produced by the RAM model for each scenario in Table 1.
- Figure 4** Illustrating the relationship between our distributed model and the constituent models used to generate its input data. The data is processed and combined to produce a risk value following the definition of (Kron, 2005). The risk value is then associated with each zipcode area in the Athens municipality and used as input to our model.
- Figure 5** Spatial variation (time $t = 0$) in four of the components used to define our risk metric.
- Figure 6** (Left): Locations of the zipcodes used in the model – a risk value is associated with each zipcode which is used to prioritise resource allocation. The risk at each zipcode evolves temporally and is updated every hour, on the hour. (Right): Locations of cooling places and metro stations in the Athens region; friendship clubs and indoor courts provide respite from the effects of the heatwave through shading and cooling.
- Figure 7** Temporally evolving snapshots of risk.
- Figure 8** Logic of the emergency worker agents.
- Figure 9** Temporally evolving snapshots of capacity produced under the logic of experiment 1.
- Figure 10** Temporal evolution of zipcode visits when one emergency worker is available. Each dot at a given time means that at least one emergency worker has visited that zipcode by the given time.
- Figure 11** Temporal evolution of zipcode visits when ten emergency workers are available. Each dot at a given time means that at least one emergency worker has visited that zipcode by the given time.
- Figure 12** Temporal evolution of zipcode visits when one hundred emergency workers are available. Each dot at a given time means that at least one emergency worker has visited that zipcode by the given time.

List of tables

Table 1	Scenarios for experimentation (Chen, et al., 2025).
Table 2	Variables used in the definition of the SHVI statistic.
Table 3	Awareness levels for high temperatures in Greece.
Table 4	Information on each of the resources implemented within the model.

Executive summary

This report presents details of the development of two distributed models implementing the *ARSINOE Dynamic Multi-System Resilience Modelling & Assessment Framework (DMRM&AF)*. The DMRM&AF, also called the ARSINOE wheel, was developed as part of D3.8. The first model is related to the Torbay area case study (CS#8) and the second to the Athens metropolitan area case study (CS#1). By distributed, we mean **soft-linked** i.e. the developed models are using outputs from other models as their inputs, but are not calling the models dynamically during run-time.

Our report begins by focussing on the Torbay case study and discusses the development of a distributed model for resource allocation. A hybrid simulation approach is implemented and the effectiveness of several resource allocation policies, such as resource sharing, are compared. As its input, the model uses output data from both the flood model CAFlood (Guidolin, et al., 2016) and a traffic model.

The second part of the report focuses on the Athens case study and again discusses the development of a distributed model for resource allocation. An agent-based approach is implemented, and a set of capacity maps are produced which represent the distribution of resources across Athens during a one-day heatwave event. The model uses output data from several models as its input. In detail, the model uses vulnerability output from a statistical model, as well as temperature and air quality output from a climate model as input.

We conclude the report by citing avenues for future development, analysis and interpretation of the models and the results.

Related Deliverables: This deliverable is directly linked to Task 3.5 (Multi-System Dynamic Modelling Framework for Resilience Assessment) and to subtask 3.5.2 (Distributed Simulation Models). This deliverable (D3.9) focusses on the technical aspects of the models implemented and is linked to D6.7 (Implementation of the ARSINOE Framework and Resilience Assessment). D6.7 contains further experimental results from the models.

1.0 Introduction

As a result of climate change, in many places extreme weather events are increasing in both severity and frequency. The importance of building an understanding of how to develop an effective emergency response strategy during such events cannot be understated and one method of doing so is to use computational modelling. It can be the case that we do not have a single model for the entire system. Rather, a set of sub-system models, each approximating a small component of the wider system, are available. The sub-models may have been developed using the extant knowledge base to different disciplines, which, in the case of ARSINOE, includes Climate Science, Water Engineering and Operations Research. In order to provide an integrated model that combines sub-models from different disciplines, the models can be linked using various approaches, for example, using the IEEE 1516 distributed simulation HLA standard and the run-time infrastructure (IEEE, 2010) for distributed and synchronised execution of the modelling artefacts, or by using the output of one model as the input to another (the approach which has been followed in the implementation of the two models presented in D3.9). Using the output of one model as the input to another is often referred to as **soft-linking**. The ARSINOE framework (refer to D6.7) has been developed based on the chaining of the models using the second approach of distributed modelling, namely, the models are executed separately (on different computers, by different teams), with the output of one model used as an input to another model.

To express the complex interdependencies between a set of models within a specific case study the ARSINOE framework, or *ARSINOE wheel* as it is also referred to, was developed in D3.8 (refer to D6.7 - Implementation of the ARSINOE Framework and Resilience Assessment). The ARSINOE wheels for the Torbay case study (CS#8) and the Athens case study (CS#1) can be found as two case study supplements that accompany D6.7.

In this report we develop two distributed (soft-linked) models; one for each aforementioned case study. In each case, the purpose of the ARSINOE wheel is to illustrate the relationship between the sub-models within the wider model; it is a conceptual model. For the Torbay case study, the main model is the Resource Allocation Model, which is also referred to as the RAM model, which is developed using hybrid simulation techniques (Brailsford, Eldabi, Kunc, Mustafee, & Osorio, 2019). The sub-models used in the resource allocation model are the CAFlood model and the SUMO traffic model. For the Athens case study, a capacity distribution model across the Athens municipal area is developed using agent-based simulation. The sub-models used within the capacity distribution model are a statistical model of vulnerability and a climate model producing temperature and air quality measurements. Full details of the development, results and further work can be found in their respective sections.

Both models developed in this report are **distributed**. The meaning of the term distributed is highly field dependent; in this report a model is said to be **distributed** if it uses the output from other models as its input. This is also known as **soft-linking** in disciplines such as climate impact modelling applications. (Wilby & Dessai, 2010)

2.0 Distributed Model 1: A Hybrid Simulation Model for Multi-Sector Resource Allocation During Flood Events

2.1 Case Study Overview

Torbay is a seaside resort town located in South Devon (UK) covering an area of approximately 62 square kilometres. Historically, Torbay has suffered from flooding on numerous occasions because of extreme rainfall events. For example, in 2014 the greatest winter rainfall since records began resulted in considerable damage due to widespread flooding. Flooding in Torbay occurs for numerous reasons including from surface water run-off and flooding of main rivers, ordinary watercourses, combined sewers and highways. Coastal overtopping events, which occur when high tides coincide with strong easterly winds, are responsible for most flooding events in Torbay. Naturally, low lying areas around the coast of Torquay, Paignton and Brixham experience the greatest impact during high tidal cycles when the capacity of the surface water outfalls discharging to coastal waters is impeded (Torbay City Council, 2014).

Impact is not restricted to residential properties; critical infrastructure, such as that used for the provision of electricity is also typically affected *directly*. More subtly, the knock-on effects suffered *indirectly* because of *direct* damage to critical infrastructure, which we refer to as cascading effects, can have considerable impact across the region. For example, an important piece of critical infrastructure during any flooding event is the region's road network, which allows the transport of resources to affected areas. In the extreme case, parts of the road network may close completely to traffic until the flood water has subsided. In a less extreme scenario, roads will not be closed entirely but the time taken for vehicles to complete their journey will still be affected; this can still have a considerable impact on provision of resources during emergencies.

As sea level is predicted to rise by over 1m in Torbay over the next 100 years, the frequency and intensity of overtopping of the sea defences will increase. Hence, to increase preparedness for future events, it is of interest to investigate strategies for the provision of resources during emergency response using a RAM model. In reality, the resources required during emergency response, such as materials, machinery, and transport, are limited and are potentially insufficient to match demand during extreme weather situations. In addition, emergency response typically involves several parties such as local government sectors, public service providers and resource requesters. Using a simulation facilitates the exploration of this complex structure and could provide guidance on emergency response strategies.

The rest of this section is structured as follows. Section 2.2 discusses the hybrid simulation methodology behind the RAM model and the data sources used in the experiments. However, since the platform we are using supports multi-methodology modelling, the model can be developed in one single model development environment and which incorporates the agent-based simulation, system dynamics and discrete-event simulation paradigms and their overarching modelling methodologies. Our model uses the output of two sub-models as input: CAFlood and the Traffic Model. Section 3.3 provides details of experimental results before the concluding remarks of Section 2.4. The details of Section 2 closely follow the methodology and results of (Chen, et al., 2025).

2.2 Methodology

2.2.1 Hybrid Simulation using agent-based simulation, discrete-event simulation and system dynamics

Due to its power and flexibility hybrid simulation (Brailsford, Eldabi, Kunc, Mustafee, & Osorio, 2019) is a technique growing in popularity. As opposed to using a single simulation technique, hybrid simulation combines two or more conventional simulation techniques within a single simulation study. Hybrid

simulation aims to offer advantages over using a single technique in isolation by leveraging the strengths and suitability of each method for its specific role within the overall model. The RAM hybrid simulation model for CS#1 comprises of sub-models developed using discrete event simulation, system dynamics and agent-based simulation (Figure 1). However, as mentioned earlier, the platform (Anylogic) which was used to develop the model provides a common modelling environment, although conceptually and methodologically they are sub-models with information and logic flows.

Discrete-Event Simulation: Discrete-event simulation models the behaviour of a system as a sequence of events that occur at discrete points in time. Between the events the state of the system is assumed to be constant. Within the RAM model, discrete-event simulation is used to model resource prioritisation, allocation and dispatch amongst competing demand from critical infrastructures. Water level dynamics involved in emergency response, such as pumping volume monitoring and water levels at critical infrastructure, are also modelled using discrete-event simulation (using the fluid library within AnyLogic).

System Dynamics: System dynamics is a general technique used to model complex systems with feedback loops. Some of the main components of a system dynamics model are stocks, flows, feedback loops and delays. System dynamics can model dynamically variable and causal relations amongst factors. Within the RAM model, it is used to model the cascading effects amongst different critical infrastructure.

Agent-based Simulation: Agent-based simulation is suited to modelling systems where, despite complex global behaviour, it is possible to clearly define the rules by which individuals act. Agent-based simulation can reveal emergent behaviour of the system that is not apparent from looking solely at individual agents. Agent-based simulation is described in more detail in Section 4.2.1. Within the RAM model we use agent-based simulation to define the logic of vehicles, pumps, materials, infrastructure and equipment.

2.2.2 Data Sources

The model uses synthetic water level data based on the ability of the CAFlood model to provide water levels at any point of interest in the output domain. The RAM model can also pull data from a traffic model, programmed in SUMO (Alvarez Lopez, et al., 2018), which estimates travel time for agents moving between critical infrastructure under flooded conditions. Further details on each of the models are included below.

Flood Data: The CAFlood model can predict flood levels in three main towns of Torbay, namely Torquay, Paignton, and Brixham following fluvial and pluvial flooding events. CAFlood can produce water levels at any location in the output domain for any sequence of timesteps. Hence, CAFlood can produce water levels at the locations of the critical infrastructure at hourly intervals. Based on this knowledge, we used synthetic data at hourly intervals for the water levels to mimic the CAFlood simulation response, but the principle remains the same. That is, CAFlood is not dynamically linked to the framework. The temporal variation in the water level at each critical infrastructure location (Figure 2) used in the model is used as an input by the other models through the *soft-linking* approach.

Traffic Model: The Traffic Model, developed in SUMO, addresses the impact of flooding on traffic (Krajewicz et al., 2012). It models traffic flow, congestion, and potential disruptions either directly caused by flooding or indirectly due to increased congestions along routes where vehicles are avoiding flooded areas. The traffic model can provide more accurate estimates of travel time for the vehicle agents compared with a model that neglects the impact of flooding on traffic completely. A version of the model that has been tested to account for this exists and is described fully by (Chen, et al., 2025). The model used for experimentation in this document *does not* directly call the SUMO model. Despite this, the existing code can be adapted to facilitate this.

2.2.3 Agent Definitions and Logic

The AnyLogic modelling environment facilitates a hybrid simulation approach by allowing each component of the RAM model to combine agent-based simulation, discrete-event simulation, and system dynamics methods. In AnyLogic, the components are referred to as agents, though this does not mean they are restricted to using only agent-based simulation. The components of the RAM model can be grouped roughly into the following categories (Figure 1):

- **Resource Hub agent:** The resource hub agent uses discrete-event simulation to implement the collective resource allocation strategy, encompassing multiple responders who own and offer their resources for emergency response. The current implementation of the RAM model has one resource hub.
- **Infrastructure agents:** Critical infrastructure agents are regarded as the local level of the RAM model. There are multiple sub-types for this agent including hospitals, electricity sub-stations and train stations. Each agent contains agent-based simulation logic allowing it to behave individually. In Anylogic, this is implemented using state-charts defining the request response according to the current flood level. A discrete-event simulation module controls the pumping process for water level control at individual critical infrastructure locations. A system dynamics module is used to model the cascading effects and relationship with other critical infrastructures, as well as for calculating the loss of service capacity for priority ranking of resource in the global level.
- **Resource agents:** Multiple instances and sub-types of the resource agent type exist. Resource agents include machines (e.g. generators, pumps and diggers), labourers (e.g. site managers), materials (e.g. sandbags), and vehicles (e.g. trucks). The agents are regarded as the local level of the RAM model. For example, the vehicle sub-type employs agent-based simulation logic to define the vehicle's delivery logic between critical infrastructure locations and the central resource hub.
- **Order agents:** These are the instances of emergency response requested by the Infrastructure agents. The order agents model the following:
 - (1) Interaction between the Resource Hub agent and instances of the Resource agent, and
 - (2) Interaction between the Resource Hub agent and instances of the Infrastructure agent.




Figure 1: Conceptual model of the RAM model with discrete-event simulation (DES), agent-based modelling (ABM) and system dynamics (SD) components .




Figure 2: Water levels at each critical infrastructure location used in the RAM model.

2.3 Results & Experiments

Our experiments compare the efficacy of the following four resource allocation strategies:

1. First come, first served.
2. Resource sharing only.

3. Priority queuing only.
4. Resource sharing and priority queuing.

Each of the allocation strategies are explained below:

First come, first served: This is the most naïve method of dealing with requests during emergency response, involving ordering the requests by arrival time and dealing with them as and when resources permit. This approach means that dealing with requests which, if dealt with immediately would have a greater benefit in the long run, could be postponed in favour of dealing with requests which will have less of an impact in the long run.

Priority Queuing: To implement the priority queuing strategy the RAM model employs a system dynamics module, using the number of people served per day as the key performance indicator, to calculate cascading effects between different critical infrastructures and create a ranking of the emergency response requests from high to low priority. This links resource priority with the cascading effect, in turn linking the resource priority to the critical infrastructure served by that resource. As a concrete example, consider a request for a resource to be delivered to a flooded power station serving n people. Suppose this power station serves one critical infrastructure, and the critical infrastructure serves a catchment of m people. This request would take priority over the same request for a power station which serves n people but does not serve any critical infrastructure. This simple example alludes to the complexity of the cascading effects when multiple requests and critical infrastructures are at play. This complexity is well captured through the system dynamics module.

Resource Sharing: Resource sharing between different agencies can address inadequacies that may exist locally by allowing resources owned by individual responders to be used by other agencies in the collective operation.

As well as comparing each allocation strategy, the experiments conducted also compare the effect of varying the total number of available resources. That is, our experiments calculate the total service capacity for each allocation strategy when:

1. A default number of resources are available.
2. Two times the default number are available.
3. Three times the default number are available.

A summary of the experimental set-up is contained in Table 1 and the results of our experiments are included in Figure 3.

From Figure 3 we see that higher resource levels, i.e. twice and three times the default number of resources, lead to an increase in total service capacity per day under each of the four resource allocation policies. However, increasing the resource level from twice to three times the default level of resources results in less of an increase than is obtained by increasing from the default number of resources to twice the number of resources. This is most obvious in under the *resource-sharing* and *resource-sharing with priority-queuing* policies, where the total service capacity per day appears to be plateauing.

The results also verify that *resource-sharing*, *priority-queuing* and the combination of the two (*resource-sharing with priority-queuing*) can substantially benefit infrastructure resilience by retaining service capacity.

For all three resource levels, the *resource-sharing*, *priority-queuing* and *resource-sharing with priority-queuing* policies provide improvements over the baseline scenario. However, the improvement varies considerably by resource level. Adopting policy *resource-sharing* or *resource-sharing with priority-queuing* can substantially benefit service capacity when the default level of resources is available, whilst adopting the policy *priority-queuing alone* cannot. This is because when emergency response resources are extremely limited, even queuing at the front position will not get the allocated resource quickly enough to realise the potential difference of capacity loss of two affected infrastructures. It is interesting to see when commonly adopting the *resource-sharing with priority-queuing*, the service capacity performs best amongst the four scenarios, implying the effect of sharing policy and the effect of priority policy work together to retain service capacity.

With twice the default level of resources available, we can see that the *resource-sharing* and *priority-queuing* policies alone can both substantially contribute to total daily service capacity. Although the *resource-sharing with priority-queuing* policy can also substantially benefit service capacity, the *resource-sharing with priority-queuing* merely performs similarly to the *resource-sharing* policy, suggesting there is little benefit from also incorporating the *priority-queuing* policy when *resource-sharing* is already being implemented. The reason behind this is that the effect of the sharing policy alone is sufficient to supply not only the requests with earlier queued positions but also those with later queued positions. This means that even the latter queued requests can receive resources shortly after those with higher priority (i.e., a negligible time gap), so the additional effect of the priority policy does not occur.

Looking at all results for three times the default level of resources, we can see all four scenarios have similar performance. This means that the sheer number of resources alone is enough to satisfy the demand and that neither the *resource-sharing* policy nor the *priority-queuing* policy is benefiting the emergency response response over a naïve first come, first serve approach. Resource availability alone is sufficient for satisfying all emergency response requests immediately, implying there could be unused resources present in the flood event. However, resources are limited in situations of practical interest.

The experimentation shows it is worth applying the *resource-sharing with priority-queuing* policy in the RAM model, as there is a high possibility that the synergy of applying the *resource-sharing* and *priority-queuing* policies will occur, particularly in real-world situations where resources are inadequately available to meet emergency response demands.

Table 1: Scenarios for experimentation (Chen, et al., 2025).

Resource Level Grouping	Scenarios	Description
Default level x1	Baseline	Baseline: Scenario without resource sharing and priority queuing (implying <i>First Come, First Served</i>)
	S	
	P	
	S & P	
Default level x2	Baseline	<i>Resource-sharing (S):</i> Resourcing sharing scenario
	S	
	P	
	S & P	
Default level x3	Baseline	<i>Priority-queuing (P):</i> Priority queuing scenario
	S	
	P	
	S & P	
		<i>Resource-sharing with Priority-queuing (S&P):</i> Resource sharing with priority queuing scenario




Figure 3: Total service capacity produced by the RAM model for each scenario in Table 1. (*S=resource-sharing; P= priority-queuing; S&P=resource-sharing with priority-queuing*)

2.4 Discussion

The distributed model for the Torbay case study successfully compared various resource allocation methods during a flooding event. As further work, we are currently in the process of developing the model further as part of a wider effort on the development of an emergency response Digital Twin (ER-DT). It is hoped that the digital twin version of the model will be enable responders to gain greater insight on effective utilisation of limited resources to enhance regional climate resilience. For this purpose, the ER-DT could incorporate models that forecast weather conditions, flood conditions and traffic conditions in real or near real-time.

As a digital twin, the forecast models would have to run efficiently alongside the simulation of the emergency response resource allocation model, with multiple models working interactively and synchronously. Moreover, the sub-models which inform the Operations Research model may need to be executed on different computers due to data access, model licence, or based on the availability of the required computational resource. To realise digital twins that incorporate various computational models, including those from Operations Research/Management Science, and execute in different simulated time-step orders, parallel and distributed simulation time synchronisation algorithms (Fujimoto, 1999) and standards like IEEE 1516 High Level Architecture (IEEE, 2010) could be used. These techniques ensure the synchronisation of simulation time and messages passed between the sub-models (Mustafee, Harper, & Viana, 2023).

3.0 Distributed Model 2: An Agent-Based Model for Resource Allocation During Heatwave Events

3.1 Case Study Overview

Since records began in 1880 the mean global surface temperature has been steadily increasing with the last ten years constituting the warmest on record. Due to the change in distribution of atmospheric conditions, resulting in greater anomalies in atmospheric conditions, the frequency and severity of extreme climate events is increasing. Heatwaves, which are characterised by high daily maximum and minimum temperatures over prolonged periods, are one type of extreme weather event and can result in devastating levels of suffering and destruction. Notable examples include the 2003 European heatwave (Stott, Stone, & Allen, 2004) and the 1995 Chicago heatwave (Semenza, et al., 1996). Athens, a diverse urban area containing densely populated built up regions, regularly experiences heatwave conditions in the summer months and is the focus of this section. As an example, in July and August 2021, Athens experienced an extreme heatwave event with temperatures reaching 43.9 degrees Celsius (Giannaros, et al., 2023).

In this study we investigate resource allocation strategies based on level of risk at the zipcode level. We describe in detail the construction of the model and perform a range of experiments to assess different strategies for resource allocation.

3.2 Methodology

3.2.1 An Overview of Agent-based simulation

Agent-based simulation is a tried and tested framework for modelling complex systems using autonomous, interacting, agents (Macal & North, 2005). The exact definition of an agent varies throughout the literature, though several agent characteristics important from a practical perspective are described in (*ibid.*) and reiterated below.

- 1) A fundamental property of agents is that they make independent decisions i.e. they act autonomously.
- 2) Agents are identifiable, discrete individuals with a set of characteristics and rules governing their behaviour and decision making. They are situated in an environment in which they can interact with other agents.
- 3) Agents are goal-directed.
- 4) Agents can be flexible and adaptive; their behaviour and decision-making rationale can evolve through their historical experiences.

Despite the simplicity of agents' local behaviour agent-based simulation can approximate complex global system behaviour, a phenomenon commonly referred to as *emergence*. Emergence can be observed in many real-world systems including ant colonies, snowflakes and the earth's climate. As systems of interest become increasingly complex, agent-based simulation merits adoption because of its ability to approximate global complexity through the aggregation of local simplicity. In addition, agent-based simulation does not require the modeler to make assumptions for the sake of analytical or computational tractability. Each of these reasons demonstrate why agent-based simulation remains an attractive approach for modelling complex systems.

Before providing details of the agents present in our model, in the next section we describe each of the data sources present in our model.

3.2.2 Data Sources (outputs of sub-models used as input for capacity model)

Temperature, air pollution, population density and vulnerability data were collected to define a risk index for zipcode areas in the Athens municipality. Figure 5 contains snapshots of each data source i.e. at a fixed point in time, whilst Figure 4 contains a schematic diagram depicting the relationship between each data source in the model, and their aggregation into the risk index. Full details of how the risk index is calculated can be found in Section 4.2.3. Importantly, the risk index is used to prioritise the allocation of resources. Details on each of the raw data sources are included below.




Figure 4: Illustrating the relationship between our distributed model and the constituent models used to generate its input data. The data is processed and combined to produce a risk value following the definition of (Kron, 2005). The risk value is then associated with each zipcode area in the Athens municipality and used as input to our model.

Temperature Data: Microclimatic simulations from the Weather Research and Forecasting (WRF) regional model (Skamarock, et al., 2019) were conducted. The WRF combines several models, including the Noah-MP land surface model (Noah-MP LSM) (He, et al., 2023) and the multilayer urban canopy model (UCM) (Ribeiro, Martilli, Zonato, & Villalba, 2021). The WRF model output consists of various meteorological parameters describing the urban climate at neighbourhood scale. Included in the model output are microclimatic parameters such as, air temperature at 2 meters (T2), surface skin temperature

(TSK), relative humidity at 2m (RH2), surface wind speed and direction, and more. For the present study, we use RH2 at hourly intervals on 28 July 2019.

Air Pollution Data: We paraphrase the description of air quality contained in Section 2.1.3 of Deliverable 0.2 (Votsi, Papangelis, Michailidou, Koutsantoni, & Gerasopoulos, 2024). Air pollution data is collected as output from the city scale chemistry transport model (CTM), EPISODE-CityChem, which is driven by emissions data from the Copernicus Atmospheric Monitoring Service (CAMS). Regional air quality predictions serve as boundary conditions. Meteorological initial and boundary conditions are provided by the Weather Research and Forecasting (WRF) regional model. A detailed description of the air quality modelling system is given in Deliverable [ARSINOE DAQ.1. NOA](#).

The high resolution CTM predicts the concentration of several pollutants, on a 100 by 100 meter grid, which is then transformed to a scalar air quality index (AQI) using the definition of the European AQI (European Environment Agency, 2019). For the current analysis, our distributed model is run over a single one-day period as a proof of concept. As such, air pollution data, i.e. the AQI values, at hourly intervals on 28 July 2019 have been selected as representative of the levels of air pollution during a period where heatwaves are more common.

Population and SHVI: The socio-economic heat vulnerability index (SHVI) (Ziliaskopoulos, Petropoulos, & Laspidou, 2024) combines eight different demographic predictors and takes a value between zero and one. A value of zero corresponds to the lowest level of vulnerability and a value of one the highest. The variables used in the definition of the SHVI are included in Table 2. In summary, SHVI combines multiple socio-economic factors to represent the vulnerability in each area as a scalar value. This is particularly useful as the risk in each area is now linked explicitly to both the demographics of that area *and* the physical information sources described above. Both SHVI and population data for 494 zipcode areas in Athens can be found in the ARSINOE Zenodo community (Ziliaskopoulos, Athens socioeconomic heat vulnerability index - features and index per region [Data set], 2023).

Table 2: Variables used in the definition of the SHVI statistic.

Variable
% Elderly
% Retired
% Living in houses built before 1980
% Living alone
% Renting
% Living in houses smaller than 60 square meters
% Unemployed
% Immigrants in developing countries




Figure 5: Spatial variation (time $t = 0$) in four of the components used to define our risk metric.

In reality, the boundaries of each zipcode are defined as the boundary of an irregular polygon. However, for the purposes of the model, agents that move must have well defined start and end *points*. Hence, within the model we use the centroids of the polygons to facilitate the point-to-point movement of agents. Zipcode areas, 494 in total, and their centroids are included in Figure 6.




Figure 6: (Left): Locations of the zipcodes used in the model – a risk value is associated with each zipcode which is used to prioritise resource allocation. The risk at each zipcode evolves temporally and is updated every hour, on the hour. (Right): Locations of cooling places and metro stations in the Athens region; friendship clubs and indoor courts provide respite from the effects of the heatwave through shading and cooling.

3.2.3 Risk

The term 'risk' is heavily overloaded and its precise definition is field dependent. For the purposes of our analysis, we adhere to the definition due to (Kron, 2005) which defines risk as the product of three terms: hazard, vulnerability and values-at-risk.

For the hazard component, we use the hazard presented by high temperatures and high degrees of air pollution. It is a natural assumption that as temperature and air pollution increase beyond a certain level an individual's level of distress will start to increase. For the vulnerability component we use the socio-economic vulnerability index (SHVI) (Ziliaskopoulos, Petropoulos, & Laspidou, 2024). Finally, for the values-at-risk component we simply use the population density. Again, it is logical assumption that risk is proportional to population density; (Kron, 2005) elaborates on this in the context of earthquake risk, making the point that a very strong earthquake which is experienced by nobody cannot present any risk. We agree with this line of reasoning and apply it here in the context of heatwave risk.

Letting AP_t^i , T_t^i , D_t^i and V_t^i denote the air-pollution, temperature, population density, and socio-economic heat vulnerability, respectively, at time t and at zipcode i , the risk at time t experienced by the population within zipcode i is thus defined by the following equation:

$$r_t^i = \frac{AP_t^i + T_t^i}{2} \times D^i \times V^i.$$

Note, population density and vulnerability are assumed to be constants which are independent of time but vary *spatially*. Hence, for the density and vulnerability terms the subscript, t , is omitted but the superscript, i , is retained.

Air-pollution, temperature and population density terms are necessarily normalised onto the unit interval to prevent quantities measured on larger scales from dominating the risk calculation; see Figure 5 and compare humidity with pollution. In detail, let $x_{\min}$ and $x_{\max}$ be the minimum and maximum values of metric¹ x at zipcode i and time t , i.e.

$$x_{\max} = \max_{i \in I, t \in T} x_t^i,$$

$$x_{\min} = \min_{i \in I, t \in T} x_t^i.$$

For a given metric x at time t and zipcode i , normalisation amounts to transforming the output onto the unit interval $[0,1]$ as follows

$$\hat{x}_t^i = \frac{x_t^i - x_{\min}}{x_{\max} - x_{\min}}.$$

Temporally evolving snapshots of the spatially varying risk over the Athens municipality during a one-day heatwave event are included in Figure 7.

¹ Note, population density at the postcode level is assumed to be a fixed quantity, so t here can be chosen arbitrarily for normalisation purposes. In addition, vulnerability is already measured on $[0,1]$ so we do not normalise the vulnerability metric.



Figure 7: Temporally evolving snapshots of risk: darker colors correspond to higher risk.

3.2.4 Heatwave Classification

There is no standardised definition of a heatwave, nor is there a universally accepted system for categorising heatwave events. One common method is to use a temperature threshold; if the temperature exceeds a certain threshold a heatwave is deemed to be occurring. This can be further developed to incorporate temporal information, such as time above a certain threshold, or indices such as humidity (Koppe, Kovats, Jendritzky, & Menne, 2004). Here, we adopt a similar, simple, approach by defining categories of heatwave using temperature thresholds, which are included in Table 2 and described below. Note, adopting this definition for heatwave categorisation means there is spatial variation in heatwave categorisation at any fixed time. That is, at any fixed time different zipcode areas can be experiencing different levels of heatwave (Stott, Stone, & Allen, 2004). Below, we include the temperature thresholds implemented within our model to characterise different levels of heatwave. These are motivated by the information on extreme temperature categorisation found at (Municipality of Athens, n.d.) and the following awareness levels used in the MeteoAlarm early warning system (EUMETNET, 2023):

Yellow (Moderate) The weather is potentially dangerous. The weather phenomena that have been forecast are not unusual, but be attentive if you intend to practice activities exposed to meteorological risks. Keep informed about the expected meteorological conditions and do not take any avoidable risks.

Orange (Severe) The weather is dangerous. Unusual meteorological phenomena have been forecast. Damage and casualties are likely to happen. Be very vigilant and keep regularly informed about the detailed expected meteorological conditions. Be aware of the risks that might be unavoidable. Follow any advice given by your authorities.

Red (Extreme) The weather is very dangerous. Exceptionally intense meteorological phenomena have been forecast. Major damage and accidents are likely, in many cases with a threat to life and limb, over a wide area. Keep frequently informed about detailed expected meteorological conditions and risks. Follow orders and any advice given by your authorities under all circumstances. Be prepared for extraordinary measures.

Table 3: Awareness levels for high temperatures in Greece.

		Awareness Levels (EUMETNET, 2023)			
			Moderate	Severe	Extreme
Source: National Meteorological Service (Greece) (Hellenic National Meteorological Service, n.d.)	Climate zone	Green	Yellow	Orange	Red
Very high temperatures	Central and southern mainland Greece	$T_{\max} < 37^{\circ}\text{C}$	$37^{\circ}\text{C} \leq T_{\max} < 41^{\circ}\text{C}$	$41^{\circ}\text{C} \leq T_{\max} < 44^{\circ}\text{C}$	$T_{\max} \geq 44^{\circ}\text{C}$

3.2.5 Resource Definitions

Our model contains a range of different resources, each modelled as an agent in the AnyLogic simulation software (Anylogic, 2017), each of which is briefly described below.

#CoolAthens Website: The #CoolAthens website raises awareness of the impact of heatwave events in the Municipality of Athens in numerous ways. It highlights the vulnerability of specific demographic groups such as the homeless and those with chronic illnesses. The website categorises heatwaves on a scale from 0-3, with 3 being the most severe. In addition, details of available resources provided by the Municipality, a heatwave survival guide, and guidance on heat stroke are available.

Circular Note: This is an informative note that is to be reposted on the websites of Municipalities, Regions, and other authorities. The note is titled "Prevention of the Impacts of High Temperatures and Heatwaves". It highlights the importance of the role the municipality plays in providing cooling places and the ability of media to disseminate information. Additionally, there is information on pathological conditions citizens may suffer because of high temperatures and guidance for dealing with them.

Metro Station Alerts: Alert messages can be provided at metro stations to inform citizens using public transport of the risks presented by heatwave conditions. The locations of the metro station can be found at <https://www.stasy.gr/en/>.

Ministry of Civil Protection Website: The ministry of civil protection website provides information on preparedness steps for reducing the impact of natural disasters, including heatwaves. General guidance for avoiding distress during heatwaves is available, as well as information specific to babies, young children and the elderly.

K.Y.A.D.A. Streetworkers: K.Y.A.D.A.'s Social Services provide on street help to the homeless population of Athens. On-site interventions are carried out by interdisciplinary teams consisting of social workers, psychologists, and a psychiatrist. Teams provide immediate and practical assistance through the provision of items such as food and clothing, and can make referrals to medical and social services as required.

Cooling Places: Cooling places are safe spaces for citizens to travel to and cool off during heatwave scenarios. We consider two types of cooling place; elder friendship clubs and indoor courts.

112 Alert Messages: In the EU 112 provides the option to send mass alerts to citizens in cases where a natural disaster or a dangerous situation which could threaten their life, health or safety, is expected or is in progress (Lelow & Marincioni, 2025). In these cases, citizens receive alerts with protection guidelines and instructions, through various communication channels.

3.2.6 Agent Logic

We now describe the agent logic for each resource; a summary of this information is contained in Table 3.

For the purposes of the model, we can think of the capacity map as a grid of discrete points, which we shall refer to as nodes, defined by the centroids of the zipcode polygons and the central hub. At the beginning of the day the capacity associated with any node is zero and as time proceeds the capacity associated with any node is increased as resources reach that node.

All but one resources are static, with only the streetworkers being modelled as dynamic agents. Streetworkers are initially located at a central hub at a latitude and longitude of (23.73, 37.98). Because the streetworkers are dynamic they move from point-to-point i.e. between nodes on the map. For

simplicity they are assumed to move between centroids of zipcode polygons and the central hub, and we assume they contribute a capacity of 1 at the node when they travel to. When they leave the zipcode to travel to another, the capacity is decreased by 1 at the zipcode they leave .

For all static agents, other than the cooling places and metro stations, the behaviour is more straightforward. Each static agent is triggered at a given level of heatwave and contributes a fixed capacity value *uniformly* over the entire municipality. Within the model, this simply corresponds to adding a fixed capacity to every zipcode node in the model.

Instead of contributing a uniform capacity across the entirety of the municipality, cooling places are assumed to contribute a capacity of 1 within a 1km radius of their centroid. Cooling places open depending on the mean heatwave level of zipcodes within a 1km radius of the cooling place. That is, if the mean heatwave category within the 1km radius is above 1 the cooling places open. Once the cool place opens, the capacity of the cool place is added to every zipcode within a 1km radius of the cooling place i.e. the set of zipcodes which were used to decide if the cooling place should open.

The capacity contribution of the metro stations is also based on a radius around zipcodes. Because we only have temperature at the zipcode locations *not* at the metro station locations, at each zipcode the product of the number of metro stations within a 1km radius of each zipcode and a capacity value of 0.7 are added to the zipcode capacity. The metro in Athens runs from 5am, so the capacity is only accounted for after 5am.

Relevant heatwave category triggers and capacities for all resources can be found in Table 4.

Table 1: Information on each of the resources implemented within the model.

Name	Type	Movement	Owner	Ray of influence	Function	Area	Capacity	Heatwave Category
#CoolAthens Website	Information	Fixed point	Municipality of Athens	No	On/Off	Entirety of Athens	0.2	0
Circular note	Information	Fixed point	Ministry of Health	Yes	On/Off	Entirety of Athens	0.2	0
Metro station alerts	Information	Fixed point	OASA	Yes	On/Off	Inside zipcode area	0.7	0
Civil Protection website	Information	Fixed point	Ministry of Civil Protection	Yes	On/Off	Entirety of Athens	0.2	0
KYADA Streetworkers	Information	Moving (point to point)	Ministry of Labour and Social Security	No	On/Off	At zipcode centroid	1	1
Cooling Places	Cooling and Shading	Fixed point	Elder Clubs: Municipality of Athens. AC places:	Yes	On/Off	1km radius around cool	1	1

			Region of Attica.			place centroid		
112 Alert Messages	Information	Fixed point	Ministry of Civil Protection	No	On/Off	Entirety of Athens	0.8	2

3.3 Results & Experiments

Each run of the model produces a capacity map, which is a set of scalar values, each associated with exactly one zipcode. The resulting capacity maps output by the model depend on several quantities, including the input data described in Section 4.2, the capacity value assigned to each agent, and the logic of the agents. We analysed the output maps generated by two different logic patterns for the streetworker agents through the following experiments:

Experiment Details

Our second experiment implements a more realistic situation whereby, assuming we have n streetworkers, the logic pattern of the streetworkers is as follows:

1. All streetworkers begin at the central hub located at a latitude and longitude of (23.73, 37.98).
2. If the time is between 7am and 7pm, for each of the n agents:
 - 2.1. The i^{th} agent moves to the zipcode with the i^{th} largest risk value.
 - 2.2. The i^{th} agent remains at the zipcode until the time, in hours, is a whole number, i.e. until the next hourly interval begins. Return to step 2.1.
3. Else: The i^{th} agent returns to central location.

The logic of an emergency worker is depicted in the flow diagram below.




Figure 8: Logic of the emergency worker agents.



Figure 9: Temporally evolving snapshots of capacity.

Streetworker visits (1)




Figure 10: Temporal evolution of zipcode visits when one emergency worker is available. Each dot at a given time means that at least one emergency worker has visited that zipcode by the given time.

Streetworker visits (10)




Figure 11: Temporal evolution of zipcode visits when ten emergency workers are available. Each dot at a given time means that at least one emergency worker has visited that zipcode by the given time.

Streetworker visits (100)




Figure 12: Temporal evolution of zipcode visits when one hundred emergency workers are available. Each dot at a given time means that at least one emergency worker has visited that zipcode by the given time.

3.4 Discussion

The model in this section helped investigate the spatio-temporal allocation of resources across Athens, guided by the risk at zipcode level. However, there remain several aspects of the model that could be improved in future work. An obvious starting point involves further optimisation and analysis of the movement of the streetworkers. The current implementation is simple, in that the i^{th} agent, from a total of n agents, moves to the zipcode with the i^{th} largest risk value and leaves on the hour. If the streetworkers were assumed to travel only to the top n zipcodes, as ranked by risk, it would be more efficient for the streetworker agents to communicate with each other and travel to the closest zipcode within the top n . From this, we could evaluate the efficiency of the streetworker movement in terms of the total distance travelled, with a shorter total distance travelled corresponding to a more efficient travel policy.

Another avenue for further development relates to the analysis of the capacity layers as each resource contributes to a given layer of the overall capacity. For example, metro station alerts can be regarded as part of an information layer and the elder friendship clubs can be regarded as part of an entirely separate cooling layer. Visualising the layers separately could aid the identification of specific regions of Athens which may be more vulnerable with respect to a given layer. This could guide the decision to provide more resources in areas which lack capacity.

4.0 Conclusion

This study explored the use of distributed (soft-linked) models for resource allocation using two case studies: flooding in Torbay (CS#8) and heatwaves in Athens (CS#1). We demonstrated the potential for improving emergency response and resource distribution using resource-sharing strategies in the Torbay case study. This motivated the conceptualisation of an Emergency Response Digital Twin (ER-DT) as a further development. By integrating real-time weather, flood and traffic forecasts, the ER-DT could provide responders with greater insight into effective resource utilisation, leveraging parallel and distributed simulation techniques for seamless model synchronisation. Further information on the Torbay case study can be found in D6.7.

The work on the Athens case study focused heavily on the details of how to combine the output of the different sub-models to construct a risk statistic used to guide spatio-temporal resource allocation at the zipcode level. We emphasised the need for enhanced movement optimisation and inter-agent communication among streetworkers. Future improvements include refining travel strategies to minimize total distance travelled and analysing capacity layers to identify vulnerabilities within the city's support infrastructure. By visualising these layers separately, decision-makers can allocate resources more effectively to underserved areas.

Together, these studies underscore the value of distributed modelling in emergency resource management. Future research could further refine these models, enhancing efficiency and adaptability to dynamic real-world challenges.

References

- Alvarez Lopez, P., Behrisch, M., Bieker-Walz, L., Erdmann, J., Flotterod, Y.-P., Hilbrich, R., . . . Wiessner, E. (2018). Microscopic Traffic Simulation using SUMO. *The 21st IEEE International Conference on Intelligent Transportation Systems*. IEEE.
- Anylogic. (2017). *Anylogic: Simulation Modeling Software Tools & Solutions for Business*. Retrieved from <https://www.anylogic.com/>
- Brailsford, S. C., Eldabi, T., Kunc, M., Mustafee, N., & Osorio, A. F. (2019). Hybrid simulation modelling in operational research: A state-of-the-art review. *European Journal of Operational Research*, 278(3), 721–737.
- Chen, O., Mustafee, N., Evans, B., Li, Q., Lewis, G., Khoury, M., . . . Savic, D. (2025). A Hybrid Simulation-Based Digital Twin for Multi-Sectoral Resource Allocation to Flooding Response. *Operational Research Society Simulation Workshop 2025 (SW25)*.
- EUMETNET. (2023). *MeteoAlarm*. Retrieved 2 2025, from EUMENET: <https://www.meteoalarm.org/en/live/>
- European Environment Agency. (2019, 1 9). *European Air Quality index*. Retrieved 2 2025, from <https://www.eea.europa.eu/en/analysis/maps-and-charts/index>
- Fujimoto, R. M. (1999). Parallel and distributed simulation. *Proceedings of the 31st conference on Winter simulation: Simulation---abridge to the future Volume 1*, (pp. 122-131).
- Giannaros, C., Agathangelidis, I., Papavasileiou, G., E., G., Kotroni, V., Lagouvardos, K., . . . Matzarakis, A. (2023). The extreme heat wave of July-August 2021 in the Athens urban area (Greece): Atmospheric and human-biometeorological analysis exploiting ultra-high resolution numerical modeling and the local climate zone. *Science of the total environment*, 857(195300).
- Guidolin, M., Chen, A. S., Ghimire, B., Keedwell, E. C., Djordjević, S., & Savić, D. A. (2016). A weighted cellular automata 2D inundation model for rapid flood analysis. *Environmental Modelling & Software*, 84, 378-394.
- He, C., Valayamkunnath, P., Barlage, M., Chen, F., Gochis, D., Cabell, R. S., . . . Ek, M. B. (2023). *The Community Noah- MP Land Surface Modeling System Technical Description Version 5.0. (Original work published 2023)*. Retrieved from <https://doi.org/10.5065/ew8g-yr95>
- Hellenic National Meteorological Service. (n.d.). Retrieved 2 2025, from <https://www.emy.gr/>
- IEEE. (2010). *IEEE 1516 - Standard for Modeling and Simulation High Level Architecture - Framework and Rules*.
- Koppe, C., Kovats, S., Jendritzky, G., & Menne, B. (2004). *Heat-waves: risks and responses* (Vol. EUR/03/5036810). World Health Organization. Regional Office for Europe.
- Kron, W. (2005). Flood risk= hazard• values• vulnerability. *Water international*, 30(1), 58-68.
- Lelow, G., & Marincioni, F. (2025). The European emergency number 112: Exploring the potential of crowd-sourced information for emergency management. *International Journal of Disaster Risk Reduction*, 105376.
- Macal, C. M., & North, M. J. (2005). Tutorial on agent-based modeling and simulation. *In Proceedings of the Winter Simulation Conference, 2005*. (pp. 14-pp). IEEE.
- Municipality of Athens. (n.d.). *#CoolAthens*. Retrieved from <https://coolathens.cityofathens.gr/en/>

- Mustafee, N., Harper, A., & Viana, J. (2023). Hybrid Models with Real-Time Data in Healthcare: A Focus on Data Synchronization and Experimentation. *Winter Simulation Conference (WSC)* (pp. 1148-1159). IEEE.
- Ribeiro, I., Martilli, A. M., Zonato, A., & Villalba, G. (2021). Highly resolved WRF-BEP/BEM simulations over Barcelona urban area with LCZ. *Atmospheric Research*, 248, 105220.
- Semenza, J. C., Rubin, C. H., Falter, K. H., Selanikio, J. D., Flanders, W. D., Howe, H. L., & Wilhelm, J. L. (1996). Heat-related deaths during the July 1995 heat wave in Chicago. *New England journal of medicine*, 335(2), 84–90.
- Skamarock, W., Klemp, J., Dudhia, J., Gill, D. O., Liu, Z., Berner, J., . . . Barker, D. &.-Y. (2019). *A Description of the Advanced Research WRF Model Version 4.1*. <https://doi.org/10.5065/1dfh-6p97> (Original work).
- Stott, P. A., Stone, D. A., & Allen, M. R. (2004). Human contribution to the European heatwave of 2003. *Nature*, 432(7017), 610-614.
- Torbay City Council. (2014, October). *Draft Extreme Weather Resilience Report 2013/2014*. Retrieved from <https://www.torbay.gov.uk/democraticservices/documents/s20790/Extreme%20Weather%20Resilience%20Report%20App1.pdf>
- Votsi, N. E., Papangelis, G., Michailidou, D. E., Koutsantoni, P., & Gerasopoulos, E. (2024). *Identification of sites for implementing NbS. ARSINOE Deliverable 0.2, H2020 grant no. 101037424*.
- Wilby, R. L., & Dessai, S. (2010). Robust adaption to climate change. *Weather*, 65(7), 180-185.
- Ziliaskopoulos, K. (2023). *Athens socioeconomic heat vulnerability index - features and index per region [Data set]*. Retrieved February 2025, from Zenodo: <https://doi.org/10.5281/zenodo.10070631>
- Ziliaskopoulos, K., Petropoulos, C., & Laspidou, C. (2024). Enhancing Sustainability: Quantifying and Mapping Vulnerability to Extreme Heat Using Socioeconomic Factors at the National, Regional and Local Levels. *Sustainability (2071-1050)*, 16(17), 7603.



This project has received funding from the European Union's Horizon H2020 innovation action programme under grant agreement 101037424.