



REMOTE SENSING, MULTI-OBJECTIVE OPTIMISATION, AND AGENT-BASED SIMULATION FOR URBAN NATURE-BASED FLOOD MITIGATION

YASHASHVI AGRAWAL (0009-0001-5916-8858)¹, MARCOS NAVARRETE PARDO(0009-0008-4692-618X)²,
ELENI NTAGKOULI³, SAMARTH PACHCHIGAR(0009-0008-6267-3784)⁴

¹ INSTITUTE FOR ADVANCED ARCHITECTURE OF CATALONIA , SPAIN

KEYWORDS	ABSTRACT
<i>Nature-based solutions Multi-objective optimisation(NSGA II) Reinforcement learning Agent-based flood simulation Remote sensing Urban flood resilience Nature-Based-Solution</i>	<i>Pluvial flooding is expected to cost roughly \$500 billion per year by 2030, fuelling interest in nature-based remedies. This study proposes a data-driven approach for detecting and assessing upstream afforestation for pluvial flood prevention in the Girona/Empordà region of Catalonia. The framework uses Google Earth Engine remote sensing, multi-criteria decision analysis, NSGA-II multi-objective optimisation, and an agent-based flood simulation with a Proximal Policy Optimisation agent at regional (21 × 21 km) and meso (1 × 1 km) scales. The pipeline proved technically feasible, and preliminary results reveal a directional improvement in simulated flood outcomes; quantifiable flood volume reduction remains a future project.</i>

1. Introduction

1.1. Context

Flooding is one of the world's most under-addressed climate concerns, with anticipated yearly damages of \$500 billion by 2030 (World Resources Institute & Deltares, 2020), despite relatively low present investment in flood prediction systems. The magnitude of this danger is not insignificant: the 2024 floods in Spain alone required around €16 billion in government-led disaster recovery costs (La Moncloa, 2024). The EU's Floods Directive (2007/60/EC) divides flood risk management into three phases: pre-disaster, during-disaster, and post-disaster. This project focuses on the pre-disaster phase, where anticipatory, nature-based intervention has the greatest potential to reduce downstream risk before it occurs.

Nature-based solutions (NbS) are increasingly being considered as both alternatives and complements to traditional hydraulic infrastructure. As noted by one of the domain experts interviewed in this research, conventional hydraulic structures are increasingly recognised as being insufficiently sustainable and adaptable to changing climatic conditions (see Section 2.1). The World Bank's catalogue of nature-based solutions for urban resilience (World Bank, 2023) identifies fourteen types of NbS, categorised according to their scale, location, approach, processes, and functions. These range from urban forests and bioretention areas to mangrove forests and salt marshes. Upland forests with deep soils can intercept and infiltrate rainfall, potentially reducing flood risk in downstream canals and rivers (Ozment et al., 2019).

1.2. Literature and research gap

Research on the role of trees in flood mitigation has expanded, although the available evidence remains uneven. Cooper et al. (2021) reviewed studies on woodland-based natural flood management in the United Kingdom, examining four woodland types: catchment, cross-slope, floodplain, and riparian. Their findings suggested that strategically planned and managed woodland can help reduce flood risk, but they noted that the available empirical data were limited. Herath et al. (2025) synthesised 50 international case studies on forest-based nature-

based solutions and identified an approximately linear relationship between forest cover and peak flow reduction. Their results indicated that a 20% increase in forest cover was associated with an approximately 10% reduction in peak flow across small, medium, and large catchments. Average reductions in peak flow varied only slightly across rainfall events with annual exceedance probabilities up to 1%. However, the effectiveness of forest cover was influenced by factors such as forest maturity, spatial distribution, and modelling assumptions. Similarly, Zhu et al. (2025), in a review of 454 natural flood management studies, identified a shift from localised interventions towards catchment-scale strategies. Their review found no strong empirical evidence that effectiveness consistently decreases with increasing catchment size or rainfall intensity, although studies covering larger catchments remain limited. At a smaller scale, Monger et al. (2024) applied a spatially distributed TOPMODEL to simulate peak flow during a 1-in-50-year storm in a 2.6 km² upland catchment in Northern England. Their simulations showed that broadleaf woodland reduced peak flow by 2.6% to 15.3%, depending on the extent and location of tree cover. Cross-slope and riparian planting produced greater reductions per 10-percentage-point increase in woodland cover, at 4.9% and 3.3%, respectively, compared with 2.7% for randomly distributed woodland.

Research on the spatial placement of nature-based flood management measures has developed alongside studies of their effectiveness. Mubeen et al. (2021) developed a suitability mapping approach for four large-scale interventions: floodplain restoration, detention basins, retention ponds, and river widening. They applied the method to the Tamnava River basin in Serbia, using flood maps to estimate the volume of water that each intervention would need to accommodate. Similarly, Liu et al. (2023) integrated remote sensing, GIS, and multi-criteria evaluation to identify potential locations for small check dams in the Mataniko catchment in Honiara, Solomon Islands. From an initial set of up to 35 candidate sites, they identified nine suitable locations for 5 m dams and fifteen for 2.5 m dams. Both studies primarily produced suitability maps or lists of potential intervention sites, focusing on identifying where measures could be implemented.

Cloud-based platforms have made spatial suitability mapping more accessible and easier to replicate across different regions. Google Earth Engine (Gorelick et al., 2017) combines a large catalogue of satellite imagery and geospatial datasets with cloud-based processing capabilities. This allows layers such as terrain, land cover, settlement patterns, and flood hazard to be assembled and analysed for different study areas without requiring users to download the datasets locally. In this framework, Google Earth Engine is used for the regional-scale analysis.

Selecting locations for nature-based interventions is commonly framed as a spatial suitability problem. GIS-based multi-criteria analysis, supported by remote sensing data, is widely used to identify suitable sites (Mubeen et al., 2021; Liu et al., 2023). However, when the selection criteria involve competing objectives, multi-objective evolutionary algorithms can generate a range of trade-off solutions instead of a single result. The Non-dominated Sorting Genetic Algorithm II (NSGA-II), introduced by Deb et al. (2002), is a widely used method for this purpose. Unlike weighted overlay approaches, which require the analyst to assign relative weights to criteria in advance, Pareto-based methods preserve the trade-offs between competing objectives and allow alternative solutions to be compared. NSGA-II has been applied to spatial planning tasks, including sustainable land-use allocation (Gao et al., 2020) and the siting of green and blue-green infrastructure (Dong et al., 2025; Rehman et al., 2024). Of particular relevance, Rehman et al. (2024) combined NSGA-II with a detailed hydrodynamic flood model to determine suitable locations for permeable surfaces, providing a close methodological precedent for the approach developed in this study.

Identifying suitable locations is only one part of the problem, as spatial suitability alone does not indicate how much flood mitigation an intervention can achieve. Agent-based platforms have been used to study pluvial flooding in urban areas, including approaches that combine sewer network models with overland flow algorithms to estimate flood extent, damage, and the effects of flood barriers (Novoa et al., 2023). Meanwhile, reinforcement learning applications in urban water management have focused primarily on controlling existing infrastructure. Mullapudi et al. (2020) used deep reinforcement learning to control stormwater systems in real time, while

Bowes et al. (2022) applied reinforcement learning to coastal urban stormwater systems to reduce flooding and improve water quality. In these applications, the agent determines how existing infrastructure should operate. The approach proposed in this study differs in its objective and timing. Using Proximal Policy Optimisation (PPO) (Schulman et al., 2017), the agent learns to select locations for flood mitigation interventions during the planning stage, before a storm occurs. Its reward is derived from the outcomes of flood simulations rather than from the real-time operation of an existing control system.

Mesa is an open-source Python framework for agent-based modelling (Kazil et al., 2020). Its models can be integrated with Gymnasium (Towers et al., 2024) to create reinforcement learning environments and trained using a Proximal Policy Optimisation (PPO) agent implemented in Stable-Baselines3 (Raffin et al., 2021). This combination forms the basis of the meso-scale stage of the proposed framework.

Taken together, these research strands address different parts of the flood mitigation problem. Studies on tree-based flood mitigation (Cooper et al., 2021; Herath et al., 2025; Zhu et al., 2025; Monger et al., 2024) investigate the effects of different woodland configurations on flood behaviour. Siting and allocation studies (Mubeen et al., 2021; Liu et al., 2023; Gao et al., 2020; Dong et al., 2025) focus on identifying suitable locations for interventions, while Rehman et al. (2024) extend this approach by evaluating candidate configurations using a detailed flood model. In contrast, reinforcement learning studies (Mullapudi et al., 2020; Bowes et al., 2022) primarily address the operation of existing infrastructure. To the best of our knowledge, no study has integrated remote sensing-based site selection, agent-based flood simulation, and reinforcement learning within a single framework. This study addresses that gap by combining these methods to support the spatial planning and flood mitigation assessment of nature-based interventions.

Despite the growing body of research, a key gap remains in the lack of data-driven, spatially explicit methods for determining where upstream afforestation should be located to maximise downstream flood mitigation and evaluate its potential effectiveness before implementation. Existing siting frameworks can identify suitable locations but often do not assess the flood mitigation benefits of individual interventions. Conversely, simulation-based approaches can estimate flood responses but may lack a systematic method for selecting where interventions should be placed.

1.3. Objective and Contribution

This research thus asks: Can we develop a data-driven method for locating and evaluating the impact of upstream afforestation on pluvial flood mitigation?

This study provides three major contributions: (1) It creates a framework that uses remote sensing, multi-criteria decision analysis, and the Non-dominated Sorting Genetic Algorithm II (NSGA-II) to discover and rank suitable sites for upstream afforestation based on a variety of site appropriateness criteria. (2) It incorporates a reinforcement learning approach, specifically Proximal Policy Optimisation (PPO), into an agent-based flood simulation to test various nature-based solution (NbS) locations and assess their effects on simulated flood volumes. (3) It combines the two components, regional-scale site selection and meso-scale flood simulation, into a single operational pipeline for the Girona/Empordà research area.

2. Methodology

The methodology follows a multi-stage computational workflow that combines geospatial data acquisition, suitability analysis, multi-objective optimization, and flood simulation with reinforcement learning. The process operates at two spatial scales: The regional scale covers a bounding box of approximately 21×21 km (nominally 20×20 km), centred on 2.9862°E , 42.0553°N for identifying suitable afforestation sites and a 1×1 km meso-scale for evaluating flood responses. Each stage uses outputs from the preceding analysis, while selected parameters

allow users to modify the suitability criteria and intervention priorities. Table 1 summarises the principal inputs, computational processes, and outputs of the framework.




Figure 1 Overview of the pipeline across the user, frontend, backend, and expert layers. The user sets the study area and NBS type, which feeds into Google Earth Engine data collection and MCDA/NSGA-II site selection at the regional scale (20×20 km), informed by domain experts. Selected patches then move to the meso scale (1×1 km), where an agent-based flood simulation with reinforcement learning evaluates the applied NBS configuration, again reviewed by experts, before results are exported.

Source(s): Authors' own elaboration, 2026.

Table 1. Overview of Computational Workflow, Inputs and Outputs

Stage	User inputs / model parameters	Computational process	Outputs
Data acquisition	Study area, selected geospatial layers	GEE-based data retrieval and preprocessing	Raster datasets
MCDA	Criteria weights for slope, soil, land cover, flood risk and urban proximity	Min-max normalisation and weighted overlay	Suitability map (0–1) Normalised
NSGA-II	Objective weights, population size, generations and spatial separation criteria	Pareto-based multi-objective optimization	Five candidate afforestation patches
Agent-Based Modelling	DEM-derived flow direction, river mask and NBS configuration	Mesa-based D8 flood routing	Simulated flood-depth distribution and volume
Reinforcement Learning	Eligible NBS patches and PPO training configuration	Gymnasium environment with PPO optimization	RL-selected NBS locations and training rewards

Source(s): Authors' own elaboration, 2026.

2.1. Qualitative grounding and expert validation

To ground the technical pipeline in practitioner and governance realities, we conducted structured interviews with six domain experts in flood risk management, sustainable water management, urban planning and disaster response. The interviewees comprised a flood-risk planner at a regional water authority, two university academics in sustainable water management and nature-based solutions implementation, a researcher in sustainable and

resilient environments, an architect with experience of a recent major flood event, and a researcher at an urban-sciences laboratory. We also consulted a local flood-disaster commission.

These interviews helped us understand Spain's institutional flood governance structure, which runs from EU policy through national ministries, regional authorities and river environmental agencies to municipalities, with research institutes working bidirectionally with the river environmental agency (Olcina et al., 2016). This structure significantly influenced our decision to organise the project around the regional water agency, presenting afforestation siting as a practical, fundable intervention rather than a purely theoretical one.

2.2. Study area and spatial scales

The study area is in the Girona/Empordà region of Catalonia, Spain. The computational pipeline uses two geographical scales: The regional scale covers a bounding box of approximately 21×21 km (nominally 20×20 km), centred on 2.9862°E , 42.0553°N . for data collecting, suitability analysis, and site selection, and a meso scale (1×1 km) for agent-based flood simulation of selected actions.

The 20×20 km regional extent was selected to cover the full city area considered in the study rather than to represent the actual upstream catchment boundary. This is appropriate because the pipeline does not derive flood hazard from a catchment-scale hydrological model. Instead, flood hazard is obtained directly from the JRC/CEMS-GLOFAS hazard map (see Section 2.3), which already incorporates the influence of the upstream catchment. Since the study does not independently simulate rainfall or runoff, delineating the actual catchment was not required for the regional suitability and site-selection stages. A fixed regional extent large enough to cover the study area was therefore sufficient for data collection, suitability analysis, and candidate-site selection. Catchment delineation would become necessary if the framework were extended to include its own rainfall-runoff or hydraulic modelling, which is identified as a direction for future work.

The 1×1 km meso-scale corresponds to the patch size used during regional-scale site selection, allowing the locations selected by NSGA-II to be transferred directly into the flood simulation without redefining the intervention areas. This scale is also small enough to support a locally detailed D8-based flood simulation for each selected patch within a reasonable computational time, while being large enough to represent the flood response to an individual afforestation intervention rather than attempting to simulate the entire basin. The second and subsequent paragraphs in sections must have 0.5 first line indentation.

All criteria for a study region centred at 2.9862°E , 42.0553°N in the Girona/Empordà area were created in Google Earth Engine (GEE). The region was generated from a 10 km buffer around this site, giving a bounding box of around 21.4×21.4 km (referred to in this study as 20×20 km). Each layer was resampled using closest neighbour to a common 30 m grid in EPSG:32630 (WGS 84 / UTM zone 30N) and exported as GeoTIFF. Resampling the coarser layers to 30 m yields no additional spatial detail. It only places all layers on a single grid. Table 2 lists the five criteria, their sources and how each was processed.

Table 2. Criteria used in the suitability analysis and optimisation.

Criterion	Dataset (GEE ID)	Native resolution	Processing
Slope	Copernicus GLO-30 DEM (COPERNICUS/DEM/GLO30)	30 m	Slope in degrees, scaled from the range 0 to 60° to 0 to 1.
Urban proximity	DLR World Settlement Footprint 2015 (DLR/WSF/WSF2015/v1)	30 m	Distance d to the nearest built-up pixel, capped at 2,500 m and inverted as $1 - \min(d, 2500) / 2500$. Built-up

			pixels were assigned a distance of 0.
Soil	OpenLandMap USDA soil texture class (OpenLandMap/SOL/SOL_TEXTURE-CLASS_USDA-TT_M/v02)	250 m	Surface (0 cm) and 30 cm classes were each converted to a suitability score (Table 2-3) and then averaged.
Land cover	CORINE Land Cover 2018 (COPERNICUS/CORINE/V20/100m/2018)	100 m	Classes were converted to a suitability score (Table 2-4).
Flood risk	JRC/CEMS GloFAS flood hazard (JRC/CEMS_GLOFAS/FloodHazard/v1)	Approximately 1 km	Maximum flood depth across return periods, divided by the maximum depth in the study region. Areas with no hazard value were set to 0.

Source(s): Authors' own elaboration based on Copernicus GLO-30 DEM, DLR World Settlement Footprint, and JRC/CEMS-GLOFAS hazard data, 2026.

The authors assigned soil and land-cover suitability scores, which are presented in Tables 3 and 4. Loam had the greatest score in soil because we thought it balances infiltration and water retention, whereas clay-rich soils infiltrate poorly and sand holds little water. Existing forest has a land cover score of zero, therefore new afforestation focuses on open area rather than wooded land.

Table 3. Soil suitability scores by USDA texture class.

Texture	Score	Texture	Score
Clay	0.3	Loam	1.0
Silty clay	0.2	Silt loam	0.9
Sandy clay	0.2	Sandy loam	0.8
Clay loam	0.7	Silt	0.6
Silty clay loam	0.6	Loamy sand	0.4

Source(s): Authors' own elaboration, 2026.

Table 4. Land-cover suitability scores by CORINE class.

Score	CORINE classes (code)
0.8	Pastures (231), land principally occupied by agriculture with natural vegetation (243), natural grassland (321)
0.7	Non-irrigated arable land (211)
0.6	Permanently irrigated land (212), complex cultivation patterns (242), moors and heathland (322)
0.5	Vineyards (221), transitional woodland-shrub (324)
0.4	Green urban areas (141)

Source(s): Authors' own elaboration, 2026.

The territory was then divided into a standard 1 km grid of $21 \times 21 = 441$ cells. Each criterion was averaged across 30 m pixels within each cell. Patches with incomplete raster coverage in

any of the five levels were deleted, leaving 359 patches to optimise. Across the 359 patches, the five patch-level values were min-max scaled from 0 to 1 using the formula

$$x_{\{\mathrm{norm}\}} = \frac{x - x_{\{\min\}}}{x_{\{\max\}} - x_{\{\min\}}}.$$

Patch positions are specified as eastings and northings in meters using the same projected coordinate system.

All five criteria are optimised; therefore, a greater result indicates a higher priority for afforestation. Steeper land scores higher because we believed it produced more runoff, therefore intercepting it offers a greater flood-mitigation benefit. Flood-prone areas receive higher scores since mitigation is more important where the risk is greater. Because the hazard layer indicates where flooding is most severe, this objective leads selection to flood-prone locations, therefore the chosen patches are viewed as priority mitigation areas rather than confirmed upstream source areas. Sites closer to settlements receive higher scores because they can protect more vulnerable assets. Higher soil and land-cover scores indicate land that is more suitable for planting and for retaining water.

2.3. Suitability mapping (MCDA)

To view the five criteria together, a weighted-sum composite was computed over the whole study region from the layers described in Section 2.3. The weights were 1:1:1:1:2 for slope, land cover, soil, urban proximity and flood risk, so flood risk carried twice the weight of any other criterion. The weights were normalised to sum to one, and the composite was rescaled to a final range of 0 to 1 (Figure 2). The weights can be changed by the user in the GEE tool to reflect different planning priorities. The composite is used for visualisation only and does not feed into the optimisation in Section 2.4.

2.4. Multi-objective optimisation (NSGA-II)

Patch candidates were ranked using NSGA-II (Deb et al., 2002), which was integrated into the DEAP library. Each individual encodes a single patch index, and its five objectives (land-cover appropriateness, slope, soil, flood risk, and urban proximity) represent the patch's normalised values. All five objectives are maximised with equal importance. Individuals are sorted according to Pareto dominance and crowding distance, thus no aggregate score drives the search. DEAP was set up with fitness weights of (1, 1, 1, 3, 1), but positive scaling does not modify the Pareto dominance order, therefore the higher flood-risk weight has no effect on whether patches are non-dominated.

Because each person encodes a single patch, the search space is limited to 359 candidates. Instead of searching over a wide space of options, the algorithm does a stochastic Pareto ranking of this collection. Crossover exchanges two individuals' indices, whereas mutation replaces one index with a random one. We utilised a population of 100 individuals, 20 generations, a crossover probability of 0.7, and a mutation probability of 0.3. Because each human has only one gene, the per-individual and per-gene mutation probabilities are same. The optimisation was repeated in five independent runs, and results are reported from the final run (run 5).

Patches in the final population were classified according to the sum of their five normalised objective values. The five highest-ranked patches were then approved in a greedy manner, with the best patch accepted first, and a patch was accepted only if it was at least 1,000 meters away from the previous patch. If fewer than five patches qualified, the space would have been decreased to 500 meters. All five selected spots are more than 1,000 meters apart (Table 5). This total is unweighted and merely acts as a ranking tool. The Pareto ranking itself has no weights. Section 3.2 details the ultimate selection.

2.5. Agent-based flood simulation and reinforcement learning

The locations identified through NSGA-II optimization were used to inform an agent-based flood simulation developed using the Mesa Python framework. The objective was to evaluate how the

spatial placement of Nature-Based Solutions (NbS), particularly upstream afforestation, could influence flood propagation and identify locations associated with improved flood mitigation. The simulation operates at a meso-scale of 1×1 km, using a cell-based spatial environment. Water movement follows the D8 flow-direction algorithm derived from the Digital Elevation Model (DEM), directing water towards neighbouring cells according to terrain gradients. A rasterized river network defines inflow locations, while a spatial buffer mask identifies areas eligible for NbS interventions.

To optimize intervention placement, the Mesa simulation was connected to a custom reinforcement learning environment (FloodNBSEnv) developed using Gymnasium. A Proximal Policy Optimisation (PPO) agent, implemented through Stable-Baselines3 with a Multi-Layer Perceptron (MLP) policy, was trained to explore alternative NbS locations. The agent's actions were represented as binary placement vectors, indicating whether an intervention was selected at a particular location. Each configuration was evaluated through the flood simulation, which returned a reward based on simulated flood volume. Configurations producing lower flood volumes received higher rewards, guiding the agent towards more effective intervention placements.

The PPO agent was trained for approximately 60,000 timesteps across two training runs. During training, it explored different spatial configurations and updated its decision policy based on the simulation feedback. The resulting outputs included selected NbS locations, simulated flood-depth maps, and training reward histories. By coupling reinforcement learning with agent-based simulation, the methodology extends spatial suitability analysis towards evaluating the potential flood-mitigation performance of different intervention locations. However, the approach remains exploratory, as a systematic comparison between baseline and intervention scenarios is still required to quantify actual reductions in simulated flood volume.

3. Implementation and Pipeline Outputs

3.1. Data pipeline and suitability mapping

We created a functioning Google Earth Engine export tool with a configurable panel for selecting from five input layers (slope, soil, landcover suitability, urban proximity, and flood risk) for a user-defined research area. The resulting MCDA suitability map, created using interactively weighted min-max normalisation, yielded a continuous 0-1 appropriateness score across the research region, with weights modifiable per criterion.

3.2. Multi-objective optimisation

The NSGA-II optimisation was performed to the five objectives mentioned in Section 2.4, resulting in a final set of five candidate patches. Instead of expressing a single optimal solution, these patches represent many trade-offs between the five objectives (Table 5; Figure 2). The highest-ranked patch (ID 181, longitude 992280, latitude 4673100, EPSG:32630) had normalised scores of 1.000 for soil suitability and flood risk, 0.242 for slope, and 0.871 for urban proximity. This demonstrates how the optimisation algorithm balanced competing site suitability factors when selecting potential locations. *Overall score is the unweighted sum of the five normalised objective values (maximum possible 5). It is a ranking aid, not a weighted fitness.*

Table 5. Final NSGA-II-selected patches from run 5, ranked by overall score.

Rank	Patch	Longitude	Latitude	Land-Cover	Slope	Soil	Flood Risk	Urban Proximity	Overall Score
1	181	992,280	4,673,100	0.087	0.242	1.000	1.000	0.871	3.200
2	231	1,004,280	4,671,100	0.832	0.022	0.914	0.432	0.899	3.098
3	287	1,004,280	4,668,100	0.875	0.061	0.534	0.775	0.844	3.089
4	126	994,280	4,676,100	0.571	0.176	0.613	0.916	0.786	3.061
5	249	1,003,280	4,670,100	0.891	0.076	0.861	0.243	0.923	2.995

Source(s): Authors' own elaboration based on NSGA-II optimisation outputs, 2026.

Land-cover suitability scores in the five selected patches ranged from 0.087 to 0.891, while urban proximity scores ranged from 0.786 to 0.923. Flood-risk scores ranged from 0.243 to 1.000, suggesting significant heterogeneity in the flood-risk objective across the alternatives considered. This variation reflects the trade-offs between the objectives, as well as the higher weighting of flood risk in the fitness function (Section 2.4).




Figure 2 Suitability Mapping

Source(s): Authors' own elaboration, 2026.

3.3. Reinforcement learning and flood simulation

The bespoke Gymnasium-based flood environment (FloodNBSEnv), powered by the Mesa agent-based model and D8 flow routing, was utilised to train a PPO agent (Stable-Baselines3, MLP policy) over about 60,000 timesteps. Over the course of two training runs, mean episode reward showed a modest improvement trend, moving from approximately -1,785 to approximately -1,770, consistent with the agent learning to reduce simulated flood volume through NBS patch placement, though reward had not fully plateaued by the end of training. The resulting spatial output plotted RL-selected NBS patches against simulated flood depth over the meso-scale research area (Figure 3).




Figure 3 PPO training and simulated flood output. Mean episode reward across two training runs (~60,000 timesteps), alongside the resulting flood depth map with RL-selected NBS patches overlaid.

Source(s): Authors' own elaboration based on PPO training and FloodNBSEnv simulation outputs, 2026.

3.4. Exploratory approaches

Parallel to the core pipeline, we investigated numerous additional approaches that were not included in the final validated system, such as HEC-RAS-based flood depth mapping for cross-

validation, an alternative urban forest GEE classifier interface, and a prototype dyke-pressure simulation. These are documented as exploratory work, not integrated pipeline components.

3.5. Discussion and Limitation

This research shows that a fully integrated, cloud-based pipeline that connects remote sensing, multi-criteria suitability analysis, multi-objective optimisation, and reinforcement-learning-driven flood simulation is technically viable for upstream afforestation site selection. Each stage of the pipeline, from raw GEE data acquisition to a trained RL policy interacting with a physically grounded flood model, was executed and yielded interpretable results.

The primary limitation of this work is that, within the project's time constraints, we were not able to complete a final, systematically tallied comparison between baseline (no intervention) and NBS-intervention flood outcomes. While the PPO training curves show a directionally consistent, improving reward trend, we do not report a specific quantified flood-volume reduction, and this figure should not be inferred from the training curves alone. Establishing this comparison rigorously, across multiple random seeds and with statistical confidence intervals, is the most important next step for this work.

Several additional directions follow directly from our own project review, including incorporating the entire nature-based solution catalogue (rather than just afforestation), automating the current manually-bridged transition between regional-scale and meso-scale data, increasing user interactivity in intervention selection, and formalising the evaluation process that connects optimisation outputs to simulation-validated outcomes.

Several additional constraints influence how the regional results should be interpreted. Flood risk is derived from a static global hazard layer, therefore the regional stage highlights priority locations but does not simulate rainfall or runoff. The coarser layers (soil, land cover, and flood hazard) were resampled to 30 m, so patch averages for these criteria correspond to the source data's natural resolution, which is around 1 km for the flood layer. Maximising hazard depth prioritises flood-prone cells, therefore the selected patches should be interpreted as priority mitigation regions rather than confirmed upstream contributing areas. The suitability scores in Tables 3 and 4 were assigned by the authors and were not tested for sensitivity, and the criterion weights in Section 2.3 are subjective. Finally, the optimisation searched only 359 candidate patches and results are reported from a single run, so the selected patches may differ in other runs.

The primary contribution of this work lies not in any individual method, as MCDA, NSGA-II, agent-based modelling, and PPO are each well-established approaches, but in their integration across two nested spatial scales. The framework connects regional-scale site-suitability assessment and selection directly with meso-scale, simulation-based evaluation of specific interventions. This addresses two common limitations in existing approaches: siting-only frameworks may identify suitable locations without evaluating how specific interventions perform, while simulation-only frameworks can evaluate flood responses without providing a systematic method for determining where interventions should be placed.

For urban and regional planners, this integration raises a practical question: how can computational methods that combine earth observation, spatial optimisation, and simulation support decisions about where nature-based interventions should be placed to improve urban flood resilience?

4. Acknowledgements

This research was carried out as the main research project of the Master's in AI for Architecture and the Built Environment at IAAC, Barcelona.

We would like to express our sincere gratitude to our guides, Oana Taut and Alex Kamenev, for guiding us throughout this process and helping us generalise our research. Without their

support, this work would not have been possible. We also thank Angelos Chronis, Areti Markopoulou, and Eleni Karafylli for providing us with the platform to carry out this research.

We would also like to thank Eva Crego (Catalan Water Agency, flood risk management and planning), Isabel Vega (UPC ETSAV, sustainable water management), Nabi Agzamov (RMIT/5th Studio, sustainable and resilient environments), Roger Sauquet (UPC ETSAV, flood risk management and NbS implementation), Artur Tudela (Batlleiroig Architects, Valencia flood history), and Mathilde Marengo (IAAC Urban Sciences Lab). We also collaborated with the Alcanar Flooding Disaster Commission. Their expertise was essential to our quantitative and qualitative analysis, and their guidance helped us define our topic, research methodology, and goals.

We are grateful to our batchmates for their active support and valuable feedback throughout the project. We would also like to thank Apoorva Tripathi, whose guidance and genuine feedback enabled us to write this paper.

Finally, we would like to thank our families and friends for their support throughout this journey.

There was no external funding for this project. Gender was not considered in the source data, which consist of geospatial, land cover, soil, and flood hazard datasets.

References

- Bowes, B. D., Wang, C., Ercan, M. B., Culver, T. B., Beling, P. A., & Goodall, J. L. (2022). Reinforcement learning-based real-time control of coastal urban stormwater systems to mitigate flooding and improve water quality. *Environmental Science: Water Research & Technology*, 8, 2065–2086. <https://doi.org/10.1039/d1ew00582k>
- Cooper, M., Patil, S., Nisbet, T., Thomas, H., Smith, A., & McDonald, M. (2021). Role of forested land for natural flood management in the UK: A review. *WIREs Water*, 8(5), e1541. <https://doi.org/10.1002/wat2.1541>
- Deb, K., Pratap, A., Agarwal, S., & Meyarivan, T. (2002). A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2), 182–197. <https://doi.org/10.1109/4235.996017>
- Dong, Y., Liu, S., Pei, X., & Wang, Y. (2025). Spatially explicit multi-objective optimization tool for green infrastructure planning based on InVEST and NSGA-II towards multifunctionality. *Land Use Policy*, 150, 107465. <https://doi.org/10.1016/j.landusepol.2024.107465>
- European Parliament and Council. (2007). Directive 2007/60/EC of the European Parliament and of the Council of 23 October 2007 on the assessment and management of flood risks. *Official Journal of the European Union*.
- Fortin, F.-A., De Rainville, F.-M., Gardner, M.-A., Parizeau, M., & Gagné, C. (2012). DEAP: Evolutionary algorithms made easy. *Journal of Machine Learning Research*, 13, 2171–2175.
- Gao, P., Wang, H., Cushman, S. A., Cheng, C., Song, C., & Ye, S. (2020). Sustainable land-use optimization using NSGA-II: Theoretical and experimental comparisons of improved algorithms. *Landscape Ecology*, 36(7), 1877–1892. <https://doi.org/10.1007/s10980-020-01051-3>
- Gorelick, N., Hancher, M., Dixon, M., Ilyushchenko, S., Thau, D., & Moore, R. (2017). Google Earth Engine: Planetary-scale geospatial analysis for everyone. *Remote Sensing of Environment*, 202, 18–27. <https://doi.org/10.1016/j.rse.2017.06.031>
- Herath, P., Croke, B., Prinsley, R., Vaze, J., & Pollino, C. (2025). A systematic review of forest cover for catchment-scale flood mitigation: A nature-based solution. *Journal of Flood Risk Management*, 18(3), e70125. <https://doi.org/10.1111/jfr3.70125>
- Kazil, J., Masad, D., & Crooks, A. (2020). Utilizing Python for agent-based modeling: The Mesa framework. In R. Thomson, H. Bisgin, C. Dancy, A. Hyder, & M. Hussain (Eds.), *Social*,

- Cultural, and Behavioral Modeling: SBP-BRiMS 2020 (Lecture Notes in Computer Science, Vol. 12268, pp. 308–317). Springer. https://doi.org/10.1007/978-3-030-61255-9_30
- La Moncloa. (2024, November 28). Council of Ministers press briefing. <https://www.lamoncloa.gob.es/lang/en/gobierno/councilministers/paginas/2024/202411128-council-press-conference.aspx>
- Liu, J., Tan, J., Zhang, S., Zhong, C., Lv, L., & Tara, A. (2023). Suitability assessment of small dams' location as nature-based solutions to reduce flood risk in Mataniko catchment, Honiara, Solomon Islands. *Sustainability*, 15(4), 3313. <https://doi.org/10.3390/su15043313>
- Monger, F., Spracklen, D. V., Kirkby, M. J., & Willis, T. (2024). Investigating the impact of woodland placement and percentage cover on flood peaks in an upland catchment using spatially distributed TOPMODEL. *Journal of Flood Risk Management*, 17(2), e12977. <https://doi.org/10.1111/jfr3.12977>
- Mubeen, A., Ruangpan, L., Vojinovic, Z., Sanchez Torrez, A., & Plavšić, J. (2021). Planning and suitability assessment of large-scale nature-based solutions for flood-risk reduction. *Water Resources Management*, 35(10), 3063–3081. <https://doi.org/10.1007/s11269-021-02848-w>
- Mullapudi, A., Lewis, M. J., Gruden, C. L., & Kerkez, B. (2020). Deep reinforcement learning for the real time control of stormwater systems. *Advances in Water Resources*, 140, 103600. <https://doi.org/10.1016/j.advwatres.2020.103600>
- Novoa, D., Reyes-Silva, J. D., Helm, B., & Krebs, P. (2023). Development and demonstration of an interactive tool in an agent-based model for assessing pluvial urban flooding. *Water*, 15(4), 696. <https://doi.org/10.3390/w15040696>
- Olcina, J. D., Saurí, D., Hernández Hernández, M., & Ribas, A. (2016). Flood policy in Spain: A review for the period 1983–2013. *Disaster Prevention and Management*, 25(1), 41–58. <https://doi.org/10.1108/dpm-05-2015-0108>
- Ozment, S., Ellison, G., & Jongman, B. (2019). Nature-based solutions for disaster risk management. World Bank Group. <https://documents1.worldbank.org/curated/en/253401551126252092/pdf/134847-NBS-for-DRM-booklet.pdf>
- Raffin, A., Hill, A., Gleave, A., Kanervisto, A., Ernestus, M., & Dormann, N. (2021). Stable-Baselines3: Reliable reinforcement learning implementations. *Journal of Machine Learning Research*, 22(268), 1–8.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal policy optimization algorithms. *arXiv*. <https://arxiv.org/abs/1707.06347>
- Towers, M., Kwiatkowski, A., Terry, J., Balis, J. U., De Cola, G., Deleu, T., Goulão, M., Kallinteris, A., Krimmel, M., KG, A., Perez-Vicente, R., Pierré, A., Schulhoff, S., Tai, J. J., Tan, H., & Younis, O. G. (2024). Gymnasium: A standard interface for reinforcement learning environments. *arXiv*. <https://arxiv.org/abs/2407.17032>
- UNDRR. (2015). The human cost of weather-related disasters: 1995–2015. United Nations Office for Disaster Risk Reduction.
- Ur Rehman, A., Glenis, V., Lewis, E., & Kilsby, C. (2024). Multi-objective optimisation framework for Blue-Green Infrastructure placement using detailed flood model. *Journal of Hydrology*, 638, 131571. <https://doi.org/10.1016/j.jhydrol.2024.131571>
- World Bank. (2023). A catalogue of nature-based solutions for urban resilience. World Bank Group.
- World Resources Institute & Deltares. (2020). Projected annual flood damage.
- Zhu, Q., Klaar, M., Willis, T., & Holden, J. (2025). A quantitative review of natural flood management research. *WIREs Water*, 12(1), e1765. <https://doi.org/10.1002/wat2.1765>