Ref. Ares(2025)1417614 - 21/02/2025

Optimal Strategies to Retain Water and Nutrients

D5.2: Post-processing & interactive visualisation of optimisation results

Authors:
Cordula Wittekind (UFZ), Michael Strauch (UFZ), Felix Witing (UFZ)

Delivery Date: 28. February 2025

This project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant agreement No. 862756.

Disclaimer

This document reflects only the author's view. The European Commission is not responsible for any use that may be made of the information it contains.

Intellectual Property Rights

Project Consortium

Document Information

Program	EU Horizon 2020 Research and Innovation Action H2020-EU.3.2.1.1 (SFS-23-2019)
Grant agreement No.	862756
Project acronym	OPTAIN
Project full name	Optimal strategies to retain and re-use water and nutrients in small agricultural catchments across different soil-climatic regions in Europe
Start of the project	September 2020
Duration	66 months
Project coordination	Prof. Dr. Martin Volk Helmholtz-Centre for Environmental Research GmbH - UFZ www.optain.eu
Deliverable	D5.2: Post-processing & interactive visualisation of optimisation results In OPTAIN's task 5.3, an R-based tool with a user interface called ParetoPick-R was developed. This tool renders the results of the multi-objective optimisation, in particular the Pareto front and its relationship with the measure implementation plans, tangible and understandable for stakeholders and researchers by providing a range of interactive visualisations. It allows to perform a clustering algorithm and assists in identifying meaningful and representative subsets of Pareto solutions while minimising information loss. To support stakeholders in finding those solutions/measure implementation plans that best match their priorities and preferences across the objectives, the tool allows to perform an Analytical Hierarchy Process.
Work package	WP5: Optimisation of NSWRM plans
Task	Task 5.3: Post-processing of optimisation results
Lead beneficiary	Helmholtz Centre for Environmental Research GmbH - UFZ
Author(s)	Cordula Wittekind (UFZ), Michael Strauch (UFZ), Felix Witing (UFZ)
Contributor(s)	Joana Eichenberger, Brigitta Szabo, Ronald Kolcsar, Mikołaj Piniewski, Mohammadreza Einikarimkandi, Matjaž Glavan, Marie Anne Eurie Forio, Natalja Čerkasova, Enrico Chiaradia, Lorenzo Sanguanini, Csilla Farkas, Christoph Schürz, Moritz Shore, Štěpán Marval
Quality check	Felix Witing (UFZ), Martin Volk (UFZ)
Planned delivery date	Month 54 (February 2025)
Actual delivery date	28/02/2025
Citation	Wittekind, C. Strauch, M. & F. Witing (2025): *Post-processing & interactive visualisation of optimisation results.* Deliverable D5.2 of the EU Horizon 2020 OPTAIN Project, Grant agreement No. 862756
Dissemination level*	PU

*PU = Public; PP = Restricted to other program participants (including the Commission Services; CO = Confidential, only for members of the consortium (including the Commission Services).

Deliverable status

Version	Date	Author(s)/Contributor(s)	Notes
0.7	05.01.2025	Cordula Wittekind (UFZ)	Development of ParetoPick-R and first draft of the deliverable report
0.8	31.01.2025	Michael Strauch (UFZ), Felix Witing (UFZ), Martin Volk (UFZ)	Revision
0.9	17.02.2025	Cordula Wittekind (UFZ)	Revision
1.0	28.02.2025	Felix Witing (UFZ)	Final

Summary

Multi-objective optimisation is a powerful approach for generating a set of Pareto optimal design alternatives that decision-makers can evaluate in order to select the most-suitable configuration. In practice, however, selecting from a large number of Pareto optimal solutions can be daunting. The objective of this report is to enable researchers and stakeholders to assess the optimisation outputs produced in OPTAINs previous Task 5.2 in a structured manner, to render the results tangible and understandable, and to maximise their use for the subsequent stakeholder consultation.

This report describes the tool ParetoPick-R, including how to run it, its data input requirements and the processes it employs. ParetoPick-R allows (1) to make the complex optimisation outputs understandable through various intuitive visualisation techniques, including for the links between the objective space and the decision space of Natural/Small Water Retention Measures (NSWRM) implementation plans. (2) It implements a methodology for reducing the high number of solutions from the previous optimisation to a manageable number while reducing information loss, and (3) allows to perform an Analytical Hierarchy Process for stakeholders to assign priorities based on pairwise preferences in a structured manner.

This report is useful for researchers and stakeholders from OPTAIN and beyond working with complex optimisation problems who want to analyse their results in a structured and meaningful way and render them actionable.

Table of Contents

Abbreviations

AHP	Analytical Hierarchy Process
COCOA	Contiguous Object COnnectivity Approach
CoMOLA	Constrained Multi-objective Optimisation of Land use Allocation
CS	Case Study
DBSCAN	Density-Based Spatial Clustering of Applications with Noise
EPI	Environmental Performance Indicator
HRU	Hydrological Response Unit
LE	Learning Environment
MARG	Multi-Actor Reference Group
MOO	Multi-Objective Optimisation
N	Nitrogen
NSWRM	Natural/Small Water Retention Measure
OPTAIN	OPTimal strategies to retAIN and re-use water and nutrients in small agricultural catchments across different soil-climatic regions in Europe
SPI	Socio-economic Performance Indicator
P	Phosphorus
PCA	Principal Component Analysis
SWAT	Soil and Water Assessment Tool
UI	User Interface
WP	Work Package

1. Introduction

A key objective of the EU H2020 project OPTAIN is the identification of optimal implementation plans for Natural/Small Water Retention Measures (NSWRMs). This includes the combination and allocation of NSWRMs at the catchment scale tailored to the characteristics and management of the OPTAIN case studies (CSs). After the most promising NSWRMs for each CS were identified through stakeholders, a set of model scenarios was run to evaluate their effectiveness, considering all sites where the implementation appeared reasonable (Piniewski et al., 2024). The modelling was done using the Soil and Water Assessment Tool SWAT+ in combination with the Contiguous Object COnnectivity Approach (COCOA) developed in this project (Schürz et al., 2022). The aim of Work Package (WP) 5 is to assess the allocation and combination of NSWRMs in an integrated way, considering environmental performance indicators (EPIs) relating to water and nutrient retention and crop yield as well as socio-economic performance indicators (SPIs) such as agricultural gross margin and implementation costs, described in detail in the deliverables D2.2 (Krzeminska and Monaco, 2022) and D4.5 (Monaco et al., 2024).

Tasks 5.1 and 5.2 of WP5 provided the case studies with the tools and support to run a multi-objective optimisation (MOO) of NSWRM implementation plans in their own catchment. In particular, each CS was tasked with optimising NSWRM plans for four different objectives; two EPIs and two SPIs (Strauch and Schürz, 2024). This optimisation process yielded a large number of multi-dimensional Pareto solutions, which are highly relevant from a research perspective, but not directly usable in the subsequent stakeholder process.

Task 5.3 outlined in this report, simplifies these MOO outputs through clustering and interactive visualisation methods. Each solution along the Pareto front is linked back to spatially-explicit management plans with detailed information on the respective NSWRM implementation. A user-friendly tool has been developed to enable researchers and stakeholders to perform these customisable clustering and interactive visualisations independently. The tool also supports the multi-actor process and stakeholder dialog in OPTAIN's Task 5.4 by providing a basis for discussion and evaluation. Moreover, it implements an Analytical Hierarchy Process (AHP) that allow to subset the large Pareto-optimal datasets based on preference-based weights for the objectives. These results will be used in the cross-case study synthesis of WP6. WP7 will embed the ParetoPick-R as a web application in the OPTAIN Learning Environment.

1.1. The need for post-processing of optimisation results

MOO is a powerful tool for exploring the full range of possible NSWRM allocations and combinations in a specific CS catchment. However, the output of MOO models can be difficult to read, to understand and to work with directly. One of the main reasons is that there are (too) many dimensions, in our case four objectives and multiple measures across a multitude of spatial units (Hydrological Response Units (HRUs) in SWAT+). This consequently results in a large set of optimal solutions.

Similar solutions, characterised by comparable objective ranges and objective combinations, can be achieved with different combinations of NSWRMs.

The relationships between measures and objectives are non-linear. They are obscured through the modelling chain simulating the complex interactions between soil, water, atmosphere and plants. These relationships could be untangled through extensive sensitivity testing, but this would go beyond the scope of our project. This means that we do not assess the effect of measure implementation/characteristics on the objective space, but do it the other way around, starting from the optimal objectives. These optima on the Pareto front have to be linked back to the combinations of measures implemented to achieve them. The relationship between decision space, the NSWRM allocation, and objective space needs to be clarified. Furthermore, trade-offs and synergies between different objectives and the implications of prioritising objectives and altering objective ranges have to be uncovered and understood. In this way, the large number of available solutions is made actionable for decision makers and stakeholders.

Interactive visualisation is required to clarify and render tangible the relationships between objectives and measure implementation plans, and between the objectives themselves. For reducing the number of Pareto-optimal solutions to a manageable number, an effective and meaningful process for eliminating and summarising Pareto-optimal solutions is needed. The structured set of solutions produced through this process requires systematic analysis to derive general as well as specific recommendations for the case studies. The Pareto front needs to be communicated and assessed with stakeholders and solutions have to be found that align best with their preferences for both objective and decision space.

1.2. Objectives

The objectives of Task 5.3 are (1) to render the Pareto front and its relationship with the measure implementation plans tangible and understandable for stakeholders and researchers. Furthermore, (2) to identify meaningful and representative subsets of Pareto solutions while minimising information loss, and (3) to support stakeholders in finding those solutions/measure implementation plans that best match their priorities and preferences across the objectives.

2. OPTAIN's post-processing concept

OPTAIN employs two main strategies for post-processing the MOO results of each CS. The first is to apply a range of visualisation tools to support the assessment, understanding and communication of the complex MOO outputs (achieving objectives 1 and 3). The second is to reduce the number of optima in a meaningful way (achieving objective 2 and 3).

Both approaches have been implemented in a tool called ParetoPick-R, which is based on R and Python (White et al., 2025b). It provides a user interface (UI) for researchers and stakeholders to analyse their respective CS's MOO outputs. In its current state, the tool requires outputs generated by the OPTAIN modelling and optimisation workflows (Piniewski et al., 2024; Schürz et al., 2022; Strauch and Schürz, 2024). This includes the fitness (or objective) values and a numerical description of the corresponding NSWRM plan for each Pareto solution as well as several SWAT+ model files.

Figure 1 provides an overview of the functionalities of ParetoPick-R and how they are connected. The following sections describe how the tool shall be used and how its functionalities and design shall achieve the objectives. Section 3 explains how to install and run ParetoPick-R. Section 4 outlines how these functionalities operate methodologically and provides examples from the OPTAIN CS Schwarzer Schöps. Section 5 provides an overview of the application across the OPTAIN case studies to date and Section 6 concludes with an outlook on the future use and development of the tool.

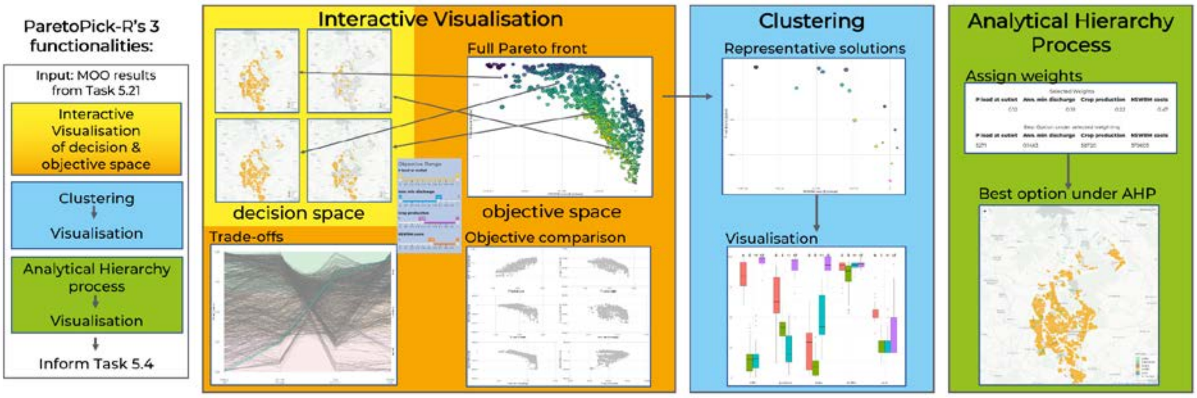

Figure 1: ParetoPick-R's functionalities and workflow

2.1. Interactive visualisation

To achieve the first objective, ParetoPick-R provides a suite of interactive visualisation tools to conduct an in-depth examination of the MOO outputs. Intuitive and interactive plotting techniques reveal patterns and trade-offs across the four-dimensional objective space (i.e. the space defined by the values of all four objectives, where each point represents a possible solution's objective values). The user can develop a comprehensive understanding of the underlying structure of the Pareto front and of the relationships among the objectives.

Using the MOO outputs produced in OPTAINs previous Task 5.2, the tool provides several types of plots with interactive elements. As described in detail in Section 4.2.3, all maps and plots are connected to four sliders where the user can limit the objective range and examine subsets of their choosing.

The first visualisation, a scatter plot (top right in the orange box in Figure 1), depicts the Pareto front across the four objectives. This allows a broad overview of the shape of the Pareto front and a general understanding of the optimisation output. The ranges of each objective and broad relationships between the objectives can be distinguished and the user might gain a first idea of potential trade-offs and synergies. A parallel axis plot of the four objectives extends the scatter plot. It provides an easy to interpret simultaneous comparison of the four objectives and further clarifies their distribution across the objective space as well as the objectives' relationships. This plot clearly unveils trade-offs through easy filtering of solutions. A collection of six scatter plots with pairwise plots of the different objectives is also provided (bottom right in the orange box in Figure 1). It allows for a direct comparison between objectives and reveals the Pareto front's actual shape and curvature from different perspectives. Potential correlations, trade-offs and synergies are clearly depicted.

The tool links the Pareto front back to each optimum's NSWRM plan, allowing for a simultaneous (3) exploration of both objective and decision space. The connection is provided by letting the user dynamically produce maps of NSWRM allocation (the decision space) by selecting optima from the Pareto front (top row in the orange box in Figure 1). This fosters a deeper understanding of the MOO process and of the effects of the previous decisions of measure allocation on the chosen objectives.

Furthermore, a dynamic frequency map is supplied. It depicts the frequency with which individual HRUs are activated in the selection made through the sliders. In this way, the importance of different measures in the catchment is emphasised and areas particularly relevant for attaining specific outcomes can be determined.

2.2. Clustering

A cluster technique was chosen to (2) reduce the number of Pareto optima in a systematic way. The aim of the clustering is to produce a smaller subset of optima while minimising information loss. It is performed based on decision space variables. In this way, some of the modelling nonlinearities can be captured and a diverse set of NSWRM plans is maintained. Similar decision spaces are considered to be more informative and practical for clustering than similar objective spaces. Stakeholders often possess an intuitive understanding of the decision space. For them it is more relevant to assess how objectives were achieved as opposed to further distinguish the optima according to their performance across the objectives.

The result of the clustering represents the objective space through fewer data points and provides a focused subset that captures objective relationships and trade-offs. The manageable number of optima reduces cognitive load for decision makers and facilitates easier comparison across measures and objectives. This also supports a better understanding of the consequences of prioritising certain objectives.

2.3. Analytical Hierarchy Process

Analytical Hierarchy Process (AHP) is a framework to assess MOO outputs in a structured manner. It simplifies and breaks down complex decision making by allowing stakeholders to compare and systematically assess conflicting objectives (Saaty, 1977). By assigning pairwise priorities, weights are allocated to each of the objectives. AHP quantifies subjective preferences into numerical values. It thereby helps stakeholders to resolve trade-offs and encourages discussion and consensus building in stakeholder groups. AHP reduces decision bias and, as a clear and traceable process, can boost stakeholder buy-in (Álvarez et al., 2013; Thungngern et al., 2017).

As part of the Multi-Actor Reference Group (MARG), Task 5.4 employs stakeholder surveys to identify the most preferred solutions at the catchment level. The post-processing results and interactive visualisations of the NSWRM plans will be used as a basis for discussion and evaluation. AHP as implemented in the tool in Task 5.3 shall support this process. It shall ensure that all opinions are considered in an equal manner and reveal the diverse, diverging and shared stakeholder preferences and priorities across the objectives (Dolan, 2008). The outcomes, solutions across the decision space (=NSWRM plans), can be discussed with regards to their feasibility and plausibility and can easily be combined with the clustering.

3. ParetoPick-R - setup and launch

ParetoPick-R is an R-based app developed with the package shiny (Chang et al., 2025). For the correlation and cluster analysis, an external Python-based executable is used (White et al., 2025b, 2025a). This section outlines the system requirements, installation, employed file structure and file inputs required to run ParetoPick-R.

3.1. Getting started

3.1.1. R installation

If R version >= 4.4.2. is not already installed on your machine please follow these steps:

1. Go to the CRAN R Project website (https://cran.r-project.org/).
2. Choose a download mirror (a server close to your location).
3. Select your operating system:
    - Windows: click Download R for Windows > base > Download R-4.4.x. for Windows.
    - Mac: click Download R for macOS and choose the version that matches your macOS.
    - Linux: follow the specific instructions for your Linux distribution (e.g., Ubuntu or Fedora).

    Install R

4. Run the file you downloaded:
    - Windows: Double-click the .exe file and follow the steps (the default options
    - work well for most users).
    - Mac: Open the .pkg file and follow the prompts to complete the installation.
    - Linux: Follow the terminal instructions provided on the CRAN website.

3.1.2. RStudio installation

Visit the RStudio website under https://posit.co/downloads/.

5. Scroll down to the RStudio Desktop section and download the free version for your operating system:
    - Windows: Download the .exe file.
    - Mac: Download the .dmg file.
    - Linux: Choose the .deb or .rpm file based on your distribution.

Install RStudio

Open the installer file:

- Windows: Double-click the .exe file and follow the instructions.
- Mac: Open the .dmg file and drag RStudio into your Applications folder.
- Linux: Use your package manager or run the file using terminal commands.

RStudio will automatically detect your R installation, so there's no extra configuration required.

3.1.3. Downloading and launching ParetoPick-R

ParetoPick-R is an app that must be run from within Rstudio. Please follow these steps to download and launch it:

Download and extract R Project:

1. Locate the zipped file on the OPTAIN cloud under WPs & Tasks > WP5 > ParetoPick-R > ParetoPick-R.zip or follow https://nc.ufz.de/s/ZExxdbtggBNJFs8?path=%2FWPs%20%26%20Tasks%2FWP5 and download it.
(*The OPTAIN cloud is only accessible for OPTAIN partners. The final version of ParetoPick-R will be made available in the OPTAIN repository on ZENODO [LINK] and on Github*)
2. Locate the zip file on your machine, right click and extract it (e.g. with 7-zip) to your preferred location

Open the project in Rstudio:

1. Launch Rstudio and in the top-left corner click File > Open File and navigate to the app folder within the ParetoPick-R folder
2. select server.R

Alternatively, open the project directly

1. Navigate to the ParetoPick-R folder
2. Double click on the ParetoPick-R.Rproj file

This will open an R script in your Rstudio viewer pane and the option "Run App" will appear at the top-right of the editor (see Figure 2).

3. Click on “Run App” to launch ParetoPick-R for the first time

Depending on your Rstudio, the app will either launch in a new window, externally in your web browser or in the viewer pane. You can switch between these settings by selecting the small down error next to “Run App”.

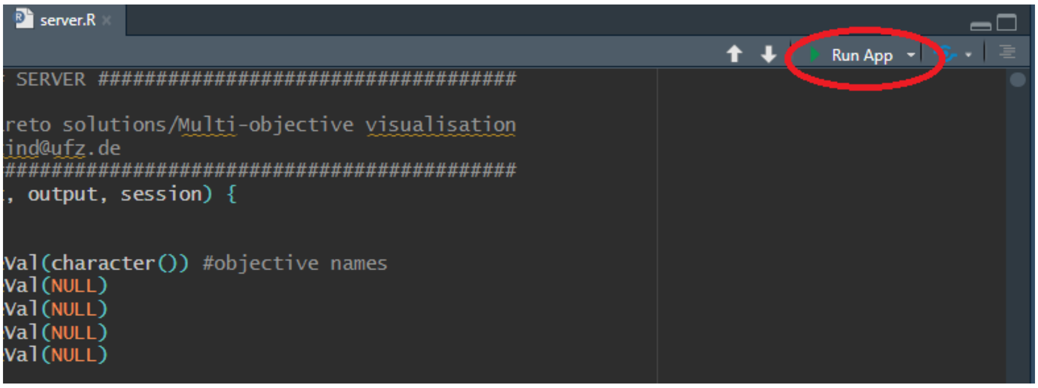

Figure 2: Running shiny-based apps in Rstudio

3.2. Data input and format requirements

Table 1 provides an overview of the data files required to run ParetoPick-R. The visualisations and interactive analytics of the “Visualisation of the Pareto Front” - tab (see section 4) require only one file, pareto_fitness.txt, and the four objective names. These are stored in object_names.RDS. To run any of the other tabs and analyses, all 10 files listed in the table need to be provided by the user.

Table 1: Overview of required files and their content. The tool copies these to the data folder.

Name	Structure	Content/comment	Source
pareto_fitness.txt	Comma (or space) delineated, four columns	The four columns represent the four objectives that were maximised during optimisation	CoMOLA_post processing.R
sq_fitness.txt	Four comma (or space) delineated values, must have same order as pareto_fitness.txt	= optional input indicating the status quo of objectives	CoMOLA_post processing.R
pareto_genomes.txt	Comma (or space) delineated list consisting of 2s and 1s	2 indicating activated and 1 indicating activated hydrological response units (HRUs) If there are x rows (=optima) in pareto_fitness.txt there should be x columns (or rows, the app understands both) in pareto_genomes.txt	CoMOLA_post processing.R
hru.con	Space separated file containing columns such as id, name, area, lat, lon, elev, hru, wst, cst, ovfl, rule, out_tot	Connection file used in the SWAT+ modelling containing details on HRU size and location	SWAT+ project folder
measure_location.csv	Comma separated table with four columns: id, name, nswrm, obj_id	File for matching individual measures and their implementation across HRUs	SWATmeasR
Shapefile consisting of four components: hru.shp hru.dbf hru.prj hru.shx		Shapefile used in SWAT+ modelling allowing the matching of HRU location and activation	SWATbuildR
rout_unit.con	Space separated file containing columns id, name, gis_id, area, lat, lon, elev, obj_id, wst, cst, ovfl, rule, out_tot, obj_typ_1, obj_id_1	Connection file used in SWAT+ modelling delineating the transport of water between HRUs, channel and aquifer	SWAT+ project folder

3.3. Folder and file structure

The ParetoPick-R app consists of six folders, as shown in the directory tree in Figure 3; the folders input, app, data, output, data_for_container and python_files.

```
.
├── app
│   ├── ui.R
│   ├── server.R
│   ├── global.R
│   └── convert_optain.R
├── python_files
│   ├── kmeans.exe
│   ├── kmedoid.exe
│   ├── correlation_matrix.exe
│   └── _internal
├── input
│   └── config.ini
├── data
├── data_for_container
│   └── config.ini (for hard reset)
└── output
```

Figure 3: ParetoPick-R's folder structure

The input folder is used for storing all data required to run the tool. These input files are internally created and regularly accessed and modified by the tool after all data files (previous section) have been provided by the user.

These files are:

- object_names.RDS: the names of the four objectives
- var_corr_par.csv: objectives and variables considered in correlation and cluster analysis
- nswrm_priorities.RDS: measures and their priority of implementation
- hru_in_optima.RDS: measure allocation across all HRUs for all optima
- all_var.RDS: all variables produced in the clustering
- pca_content.RDS: variables considered in the clustering after highly correlated variables have been removed from all_var
- config.ini: used for communicating with the external Python processes
- buffers.RDS: names of measures that require a buffer to improve their visibility in maps

The folder "data for container" contains the default configuration file, called config.ini, which is used by the external python executables. During a reset of the app, this file is used to restore the config.ini in the input folder.

All files supplied through the UI are stored in the data folder, these are the outputs of the previous optimisation process carried out in OPTAINs Task 5.2, using common optimisation protocol of Task 5.1 (Strauch and Schürz, 2024), see Table 1 for details.

The output folder is where all the files produced during the correlation and cluster analysis are written to. When selecting to save specific optima, these will be written to selected_optima.csv, which is also stored in this folder.

The folder python_files contains the three Python-based executables, correlation_matrix.exe, kmeans.exe and kmedoid.exe. Furthermore, it contains a folder called _internal required to run these executables by creating a temporary Python environment including all necessary dependencies.

The app folder contains the five R scripts for running the app: ui.R, app.R, server.R, global.R, and convert_optain.R. Each script serves a specific purpose in the app's architecture:

1. ui.R: This script defines the UI of the app. It layouts the app's structure, including input controls for data preparation and selection, clustering parameters, and visualisation options. It also defines output areas for displaying plots, tables, and clustering results.
2. server.R: This is the central code which contains the server-side logic of the app. It handles user inputs, processes data, and updates outputs. It uses reactive expressions to efficiently manage data flow, calls functions from functions.R and defines its own to create dynamic visualisations and tables.
3. functions.R: This script defines all custom functions used throughout the app. These are mainly formatting, data manipulation and custom plotting functions, but also several functions for adapting config.ini to control the external Python processes. By centralising the function definitions here, the code remains easier to maintain.
4. global.R: This short script sets up global paths and app settings. It installs packages and defines constants such as file paths, default parameters, and any configuration options that need to be accessible across the entire app. It's kept concise to focus on app-wide settings.
5. convert_optain.R: This script handles all data preparation and is needed to use the OPTAIN-specific project data in PartoPick-R. It reads the supplied files and prepares the input data for the clustering analysis, including the selection and preprocessing of variables.

This structure separates concerns effectively, making the app modular and easier to develop and maintain. The convert_optain.R script ensures data consistency for the clustering across OPTAIN case studies.

4. ParetoPick-R - functionalities

4.1. Main components and structure

The app's main interface consists of a sidebar menu on the left and a main panel on the right. The sidebar contains all tabs and a glossary defining abbreviations commonly used throughout the app. Next to the ParetoPick-R slogan over the sidebar menu, a hamburger menu (☰) allows to toggle the sidebar on and off. Depending on the screen width, the sidebar covers 20% of the total display width. Turning the sidebar off can be helpful for improving the visibility of plots and tables, especially on smaller screens. Some of the tabs have an additional sidebar menu containing explanations, controls and sliders.

ParetoPick-R's eight tabs, six of which are always displayed in the sidebar menu, are: Introduction, Data Preparation, Visualising the Pareto Front, Configure Clustering, Correlation Analysis, PCA & kmeans/kmedoids, Cluster Analysis and AHP (Figure 4).

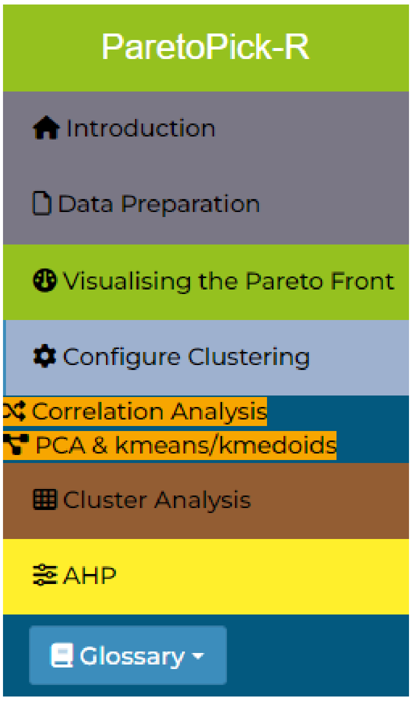

Figure 4: ParetoPick-R's main navigation sidebar

The two tabs, Correlation Analysis and PCA & kmeans/kmedoids, are only visible and accessible when the user decides to perform a manual clustering and clicks on "show cluster tabs" in the Configure Clustering tab.

4.2. ParetoPick-R tabs and their functionalities

Accompanied by example figures from the German CS Schwarzer Schöps, this section outlines how interactive visualisation, clustering and AHP are implemented across the app's tabs.

4.2.1. Introduction Tab

The introduction tab opens when ParetoPick-R is launched. It contains only text and briefly summarises the clustering and AHP methodologies and provides a general overview of the other tabs' functionalities. It outlines the main settings available and what each requires the user to do.

4.2.2. Data Preparation Tab

This tab walks the user through the required data files described in Section 3.2 and lets them upload these files through a window that opens when clicking "Browse" below the file's name. The UI allows users to select files only if they are in the correct data format.

The tab has two sections. The first section is for providing the main MOO output, pareto_fitness.txt, the objective names and, optionally, the objectives' units. This represents all files and information needed to plot the interactive Pareto front in the next tab. For procedural reasons, many parts of the app do not work without these files, each of the inputs in this section require the user to click "save" separately.

The files in the second section have to be provided if further analyses and map plotting shall be performed. After uploading all files correctly, clicking "Check Files" returns "All Files found". If there are files missing, clicking "Check Files" will return a list of all missing files. Only if all required files have been found, the option "Run Prep" will become available. As shown in Figure 5, clicking "Run Prep" calls the external data preparation script convert_optain.R and its outputs are printed to the UI. This external script produces the decision space variables required for the subsequent clustering and writes them to var_corr_par.csv in the input folder. These variables are described in section 4.2.4 (Configure Clustering Tab).

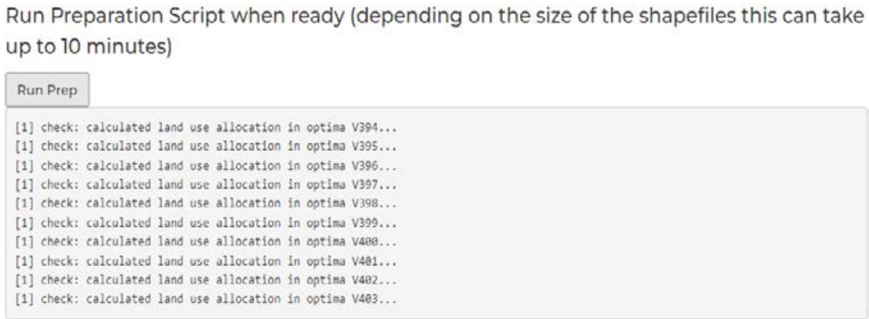

Figure 5: Running the data preparation. Screenshot of ParetoPick-R.

The Data Preparation tab furthermore contains the option to perform a hard reset to return the app to its original empty state. This hard reset deletes all previously uploaded and produced files from the input, output and data folders and copies an empty config.ini file to the input folder. As outlined below, this config.ini file is needed for calling the external python executables.

4.2.3. Visualisation of the Pareto Front Tab

The main tab for visualising and developing an understanding of the Pareto front has a sidebar on the left, containing sliders of the four objectives as shown in Figure 6. For easier comparison across different scales and for ensuring shiny's performance, the objectives are scaled to values between 0 and 1 with 1 aligning with the best MOO outcome of the respective objective and 0 with the worst. Figure 6 hence represents a selection for high performance in Phosphorus (P) load (minimised) and annual minimum discharge (maximised) as well as costs (minimised). The selection accepts optima with below average performance (< 0.5) in crop production.

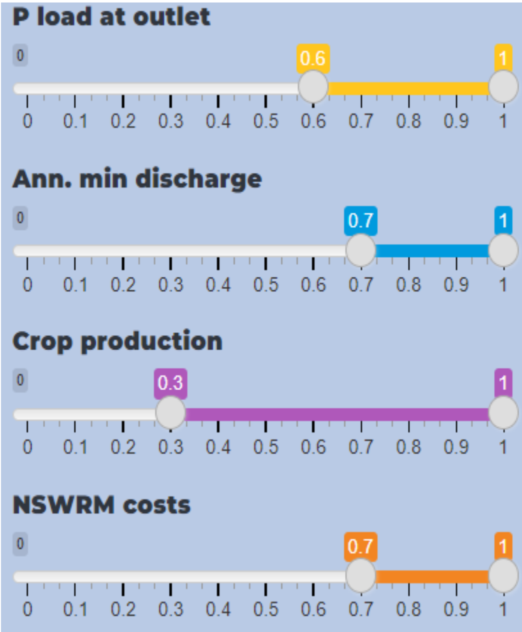

Figure 6: Objective sliders in the Visualisation of the Pareto Front tab. Example from CS 1 Schwarzer Schöps.

Below the sliders, a map of the catchment showing the frequency of measure implementation is provided. It depicts with which frequency individual HRUs are part of the remaining Pareto front as selected in the sliders. For simplicity only three different hues are provided to distinguish low, medium and high frequency. Figure 7 shows an example of this map for the selection made in Figure 6. From the map it can be seen that grass slopes (grassed waterways) play an important role for achieving these optima and are implemented with a high frequency, as are ponds and buffers. Hedges appear to be less relevant, likely because they have comparatively high implementation costs. Low till is frequently implemented across these optima and some areas, such as the fields in the west of the catchment, appear to be more important than others for achieving these objective ranges.

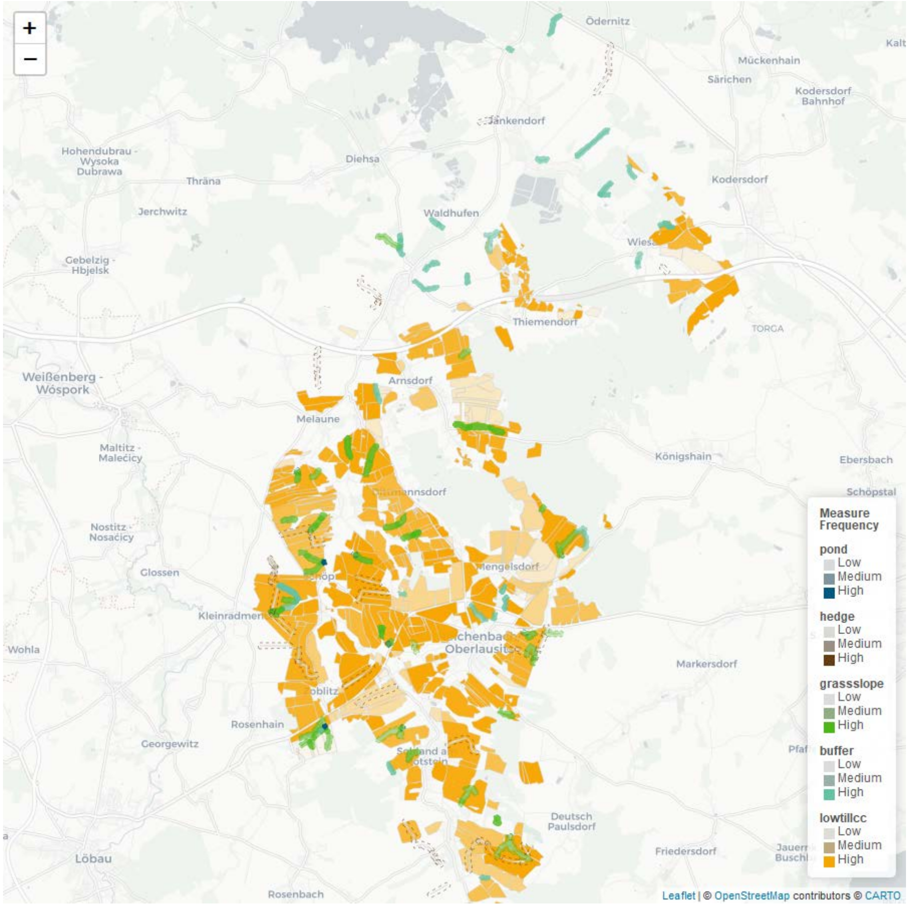

Figure 7: Summary of all measure implementation plans. Frequency with which individual HRUs are activated in the slider selection of Figure 6. Example from CS 1 Schwarzer Schöps.

On first launch, the main panel on the right contains three plots and four tables. On interacting with the sliders and plots, another map and two more tables are shown as outlined below.

The first plot of this tab is a scatter plot of the Pareto front, shown in Figure 8. The user can select which objectives shall define which setting for the two axes, for colour and for shape. On first launch, the objective with the largest absolute range is shown on the x axis and the objective with the second largest absolute range is shown on the y axis. If the objectives selected for x and/or y axis have been optimised for negative values, i.e. if the goal of optimisation was to minimise their values, the user can reverse the x and y axes. Furthermore, the status quo and the units can be added to the plot by selecting and deselecting checkboxes.

Figure 8 shows that there is a trade-off between the P loads and implementation costs in the German CS; in optima with lower P loads, implementation costs tend to be higher. However, this trade-off has a convex shape, indicating a high potential for 'good' compromise solutions. For example, it may be possible to reduce average annual P loads to less than 6000 kg with very low NSWRM costs or even savings (at catchment scale). These solutions (P loads < 6000 kg/year and costs around 0) in turn have a trade-off with the other two objectives, minimum annual runoff and crop production. However, the overall magnitude of change and therefore the risk of loss for these two objectives is comparatively small (<5%).

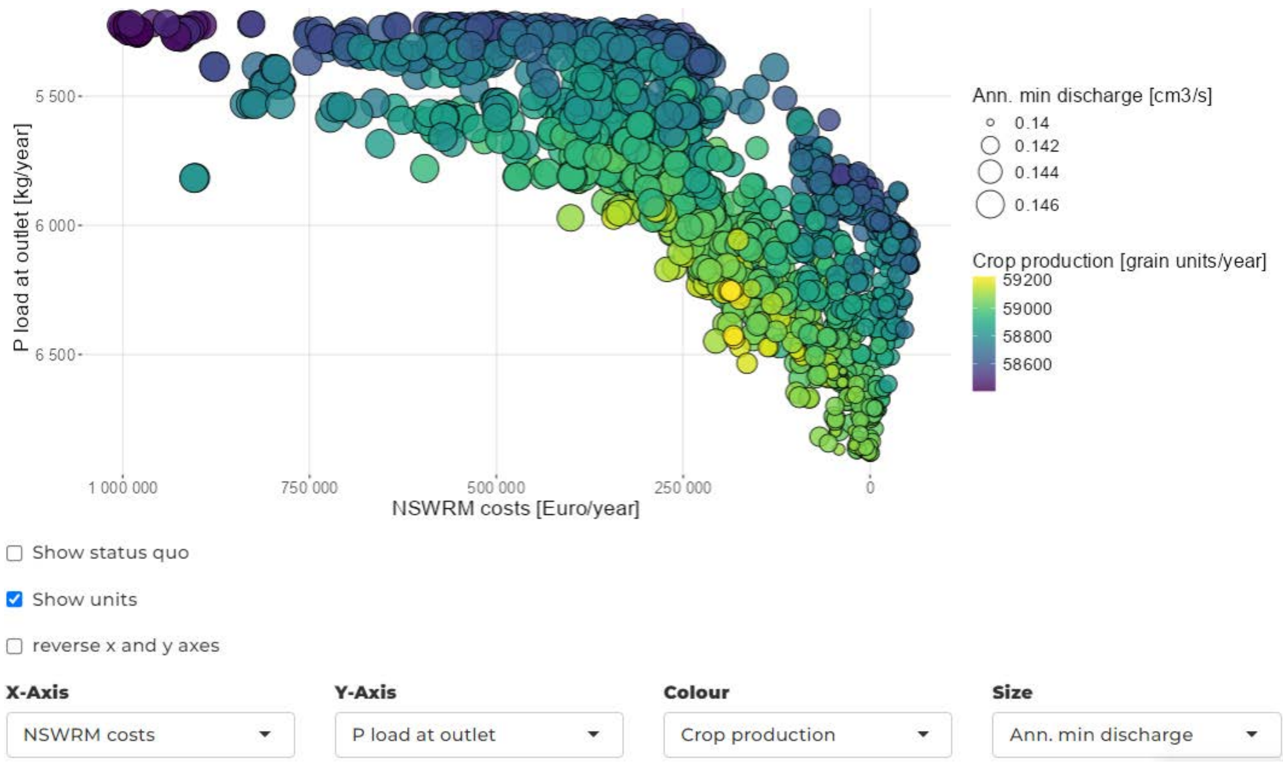

Figure 8: Scatter plot of the full Pareto front and set of visualisation options. Example from CS 1 Schwarzer Schöps.

To bridge the gap between the objective space depicted through the Pareto front and the decision space (NSWRM implementation), the user can click on individual optima and plot the map of the measure allocation in the catchment required to attain this optimum. The users have the option to zoom and pan across the map. Comparing the NSWRM plans of several optima can potentially reveal broad relationships between decision and objective space. Figure 9 shows two example maps. One optimum, number 1413, was selected for the optimisation goals of high annual minimum discharge (0.147 cm3/s) and low phosphorus (P) loads (5,268 kg/year). The other optimum, number 1400, is an example for high crop production (59,226 grain units/year). Achieving good results for both environmental objectives requires the implementation of a large number of different measures across the catchment. High crop production on the other hand entails the implementation of only low till management and a few buffers. Accordingly, high crop production can be realised with lower implementation costs. It should be noted that both optima are located on the fringes of the Pareto front. Good performance across both environmental as well as in the economic objectives is realised through other optima with few trade-offs compared to the extrema chosen here.

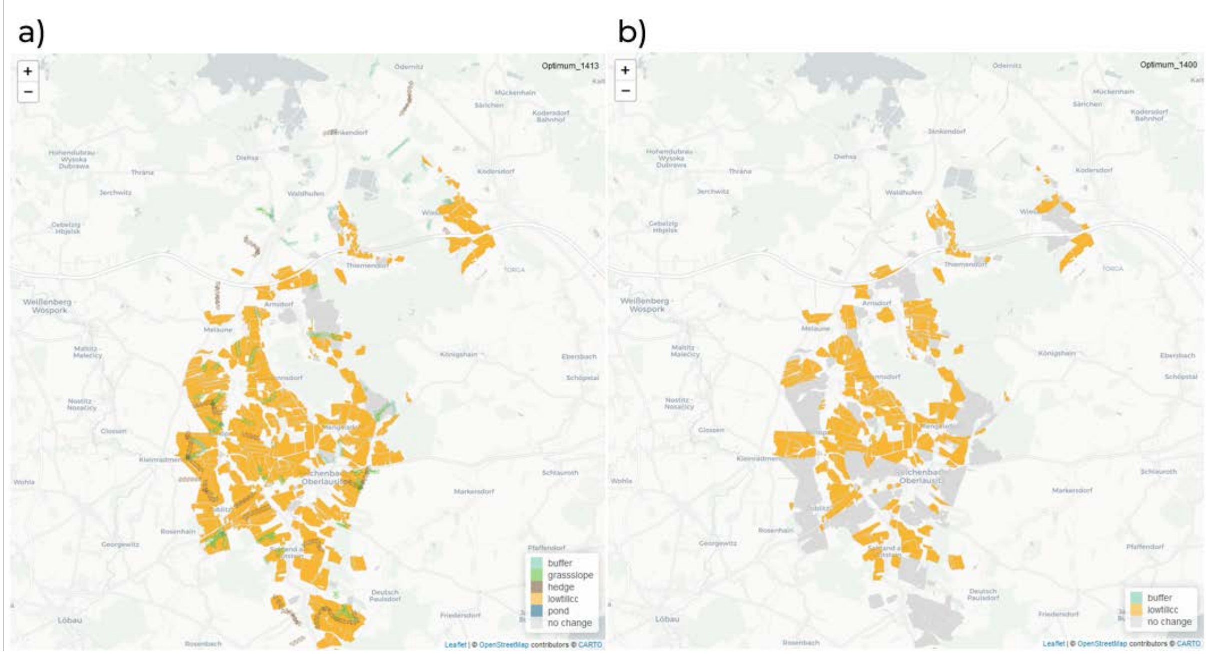

Figure 9: Measure implementation plans; (a) optimum 1413 - high annual minimum flow and low P loads and (b) optimum 1400 - high crop production. Example from CS 1 - Schwarzer Schöps.

The Pareto front plot is accompanied by two tables. The first, an example of which is given in Table 2, provides an overview of the number of implemented measure types (hedges, ponds, etc.) across the selection made in the sliders (first value) compared to the total number of measures available (second value). If an optimum is selected in the Pareto plot, the second table is visible and provides the number of individual measure types implemented in the NSWRM plan of this optimum. An example of this table is provided in Table 3. In line with the corresponding map of this optimum (Figure 9(a)), which shows a high coverage of measures across the catchment, almost all measures available for implementation in the German CS (Table 2) have been used to achieve this optimum. 27 out of 30 possible buffers are used and 26 out of 28 hedges have been implemented. However, significantly fewer ponds are used than possible. The number of hedges is higher than the selection referenced in Table 2 (26 vs. 13) because this optimum is not included within the slider selection shown in Figure 6.

Table 2: Number of implemented versus total number of measure types. The selection aligns with the sliders shown in Figure 6. Example from CS 1 Schwarzer Schöps.

buffer	grassslope	hedge	lowtillcc	pond
30 / 30	34 / 34	13 / 28	200 / 201	8 / 9

Table 3: Number of individual measures implemented under optimum 1413 as mapped in Figure 9(a). Example from CS 1 Schwarzer Schöps.

buffer	grassslope	hedge	lowtillcc	pond
27	31	26	182	2

The next two tables below provide the scaled and absolute/actual objective ranges as selected in the sliders. The latter, as shown in Table 4, also includes the percentage difference between (currently) selected values and the total minima and maxima of the Pareto front. This shall support an understanding of the scaling to between 0 and 1 and inform on the accepted trade-offs across objectives, including indirect losses, when moving the sliders. In the example of the German CS, the selection of 0.6 (P load), 0.7 (minimum flow), 0.3 (crop production) and 0.7 (NSWRM costs) also leads to a reduction of the best attainable objective values, most strongly seen in crop production. The reason is that many optima with high crop production tend to perform poorly on environmental objectives, particularly P load, leading to their exclusion.

Table 4: Selected Objective Ranges (absolute) under the slider settings of Figure 6. Example from CS 1 Schwarzer Schöps.

	P load at outlet (kg/year)	**Ann. min discharge (cm3/s)**	**Crop production (grain units/year)**	**NSWRM costs (Euro/year)**
best	5272 (-3.2%)	0.1464 (-7.44%)	59008 (-26.6%)	-93475 (-14.06%)
worst	5878 (+60.37%)	0.1448 (+70.39%)	58666 (+31.64%)	-261198 (+70.07%)

The parallel axis plot below these tables allows users to compare solutions in a different way, as shown in Figure 10. The objectives are scaled to between 0 and 1 and divided into three ranges, distinguishing worst, medium and best. There is also an option to show the status quo. Despite the strict selection for costs, focusing only on those optima beyond (cheaper) than 0.7 (= 26,1782 €), there are a considerable number of optima that perform well across the two environmental objectives.

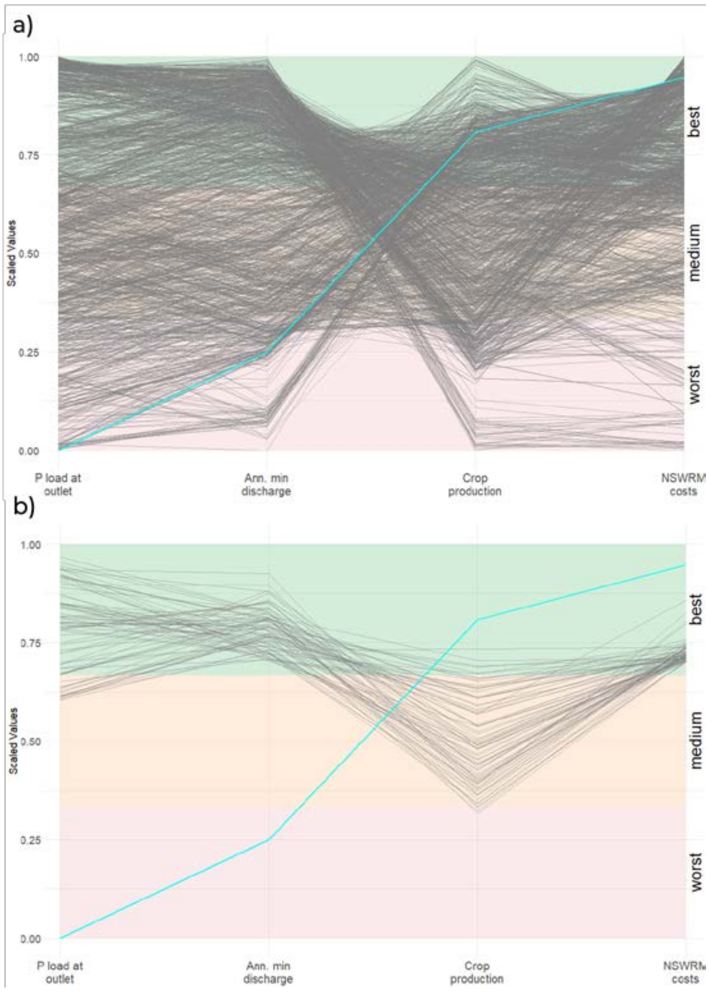

Figure 10: Parallel axis plot of the full Pareto front (a) and the Pareto plot under the selection made in Figure 6 (b) with status quo in cyan. Example from CS 1 Schwarzer Schöps.

Of the two tables at the bottom of this tab, the left one is only displayed when a selection is made in the Parallel axis plot and provides the respective absolute values of the objectives in this optimum. The table on the right, an example of which is provided in Table 5, is not connected to any interactive elements and does not change. It shows the absolute maximum and minimum for the four objective ranges, providing a reference frame for the scaled and actual objective values that have been optimised. Furthermore, as with the NSWRM implementation costs for the German CS, it emphasises where objectives have been optimised across both negative and positive ranges by adding minus and plus signs to their minima and maxima.

Table 5: Maximum Objective Ranges (absolute). Example from CS 1 Schwarzer Schöps.

	P load at outlet (kg/year)	**Ann. min discharge (cm3/s)**	**Crop production (grain units/year)**	**NSWRM costs (Euro/year)**
best	5219	0.1469	59226	+55171
worst	6882	0.1399	58406	-1001882

The last plot of this tab is a multi-panel plot. Six scatter plots illustrate pairwise objective relationships. This shall delineate the Pareto front's shape in more detail. The option to show the status quo is also provided. For the German CS (Figure 11), the plot reveals a win-win relationship between the two environmental objectives (minimum discharge and P load), while trade-offs are visible for each combination of environmental and economic objectives (e.g. crop production and P load). The two economic objectives are generally positively correlated. Nevertheless, the respective optima are relatively far apart, probably because some greening measures (grassed waterways and grassed riparian buffers) reduce the operating costs of land management (resulting in negative NSWRM costs, i.e. net savings), while at the same time reducing the value of crop production.

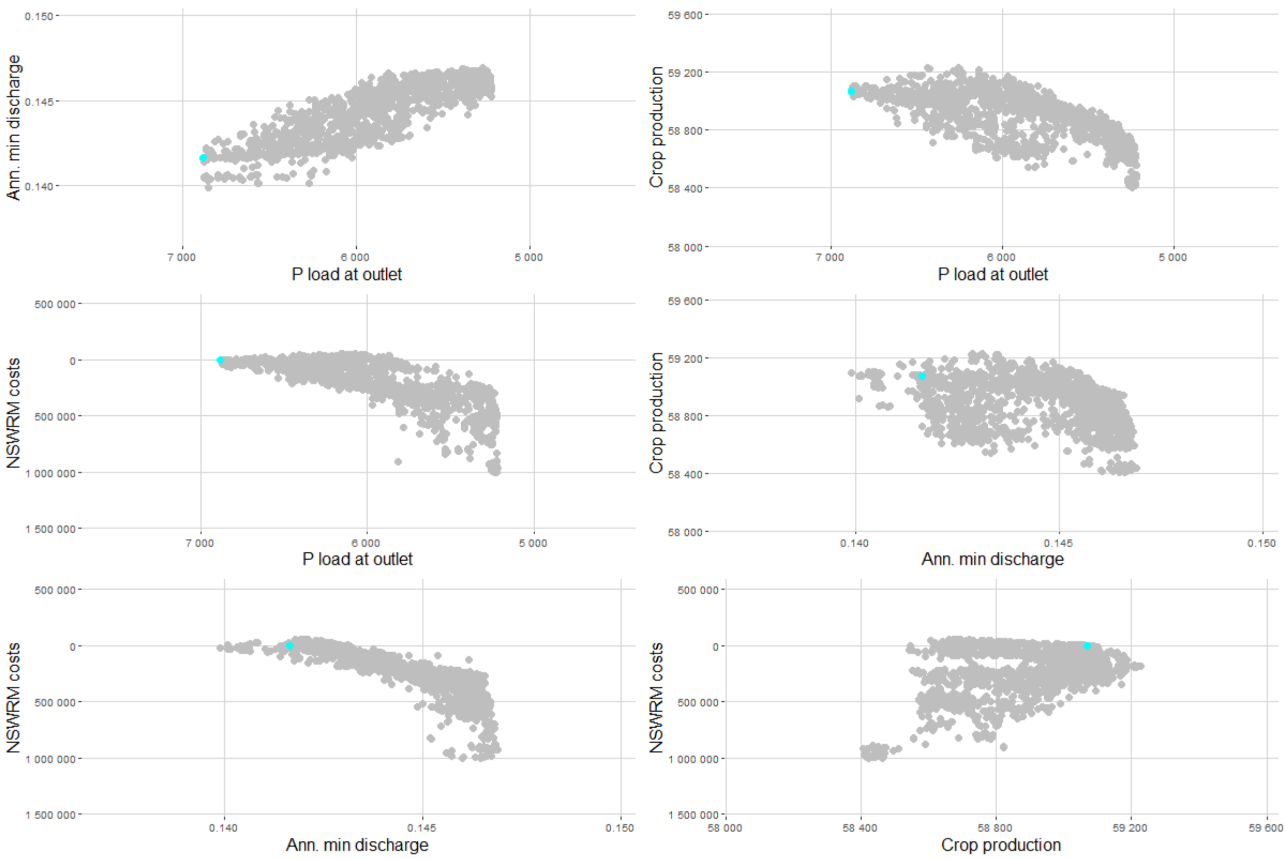

Figure 11: Pairwise scatter plot of the four objectives with status quo in cyan. Example from CS 1 Schwarzer Schöps.

4.2.4. Configure Clustering Tab

This tab is for controlling and deciding on the overall cluster process. When opening the Configure Clustering tab a simple layout depicted in Figure 12 appears. The user is requested to decide if they want to limit the objectives prior to clustering. On clicking "Yes", sliders like in the previous tab appear and if the user decides to limit the objective ranges, a reduced set of pareto options is considered in the subsequent clustering. Two options are provided below; the first is to perform the clustering with default settings, the second allows the user to jump to the

Correlation Analysis and PCA & kmeans/kmedoids tabs where they can set, refine and test a range of settings for the correlation and cluster processes.

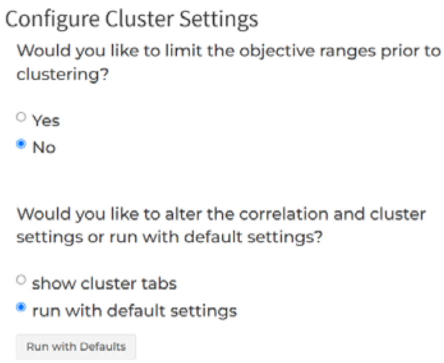

Figure 12: The Configure Clustering tab on first opening

The aim of providing a default cluster option was to allow all case studies to receive reasonable results without having to optimise the settings manually. The default cluster run uses preset values for the number of clusters, accepted degree of correlation and number of principal components (see also Table 5). Depending on the MOO outputs the cluster performance will vary. The default settings produced reasonable results for seven example datasets tested, six of which stemmed from OPTAIN CSs. It is however recommended to perform at least one manual cluster run to gain an understanding of the process.

Choosing the default cluster option through clicking on "Run with Defaults" calls the Python executables. If the scripts are executed successfully, their outputs, three figures, open in separate windows one after the other. For each of them, the user has to decide to either save or close them. The user can then continue to the Cluster Analysis tab.

Performing a "manual" clustering through the cluster tabs allows to test and optimise a range of settings as outlined in Section 4.2.6. This approach therefore likely leads to better cluster outputs more adapted to the individual CS, its decision space and Pareto front. The method chosen for the clustering is a diversity-based *a posteriori* Pareto pruning method (Petchrompo et al., 2022) that aims to provide a balanced overview of possibilities, minimising information loss. As outlined in Section 2, the decision space was deemed more interesting for the clustering, as the strategies for attaining objectives are more relevant for stakeholders than a detailed distinction of optima along the objectives.

The following variables describing the decision space were chosen as basis for the pruning ("measure" refers here to hedge, pond, terrace etc.; "measure type" distinguishes land use, management and structural types of measures):

- Each measure's share in the total catchment area considered for implementation (share_con).
- The median spatial autocorrelation between HRUs allocated to individual measures, Moran's I (Moran, 1950).
- The median fraction of water of each activated HRU of a measure that is routed directly into the channel (channel_frac).

- The ratio between structural and management measure types (linE).
- The share of land use measures in the total catchment area considered for implementation (lu_share).

Moran's I, channel_frac and share_con are calculated separately for each implemented measure. Since the number and type of measures vary across case studies, the number of variables considered in the correlation analysis and subsequent clustering therefore also varies in different case studies.

The method applied for clustering the Pareto front is outlined in detail by White et al. (2025a, 2025b) and has been implemented in a Python-based executable that is linked to and provided with ParetoPick-R. The process relies on the two commonly applied cluster methods kmeans and kmedoid (Suarna et al., 2021). Prior to clustering, a Principal Component Analysis (PCA) and a correlation analysis are performed. When performing the clustering through the cluster tabs, cluster settings are optimised through iterative testing of different user-defined settings as explained in Section 4.2.6 (PCA & kmeans/kmedoids Tab). The best cluster setting is decided based on Silhouette scores (Rousseeuw, 1987). The result is a number of Pareto optima representative for each identified cluster.

The three figures produced by the python executables that open in three windows one after the other are; a table with an overview of the percentile distribution across the objective ranges for 0-33%, 33% - 66% and 66% - 100% for each of the clusters, the Pareto front plot with the remaining solutions that are representative for individual clusters and a violin plot depicting the distribution across the objectives for each cluster.

4.2.5. Correlation Analysis Tab

The first of two tabs in the manual clustering process allows the user to decide which correlated variables to remove prior to the main cluster analysis. Although not strictly necessary, removing correlated variables prevents over-emphasis on certain features and improves the interpretability of the subsequent clustering. The sidebar on the left contains the interactive process divided into five steps, as shown in Figure 13.

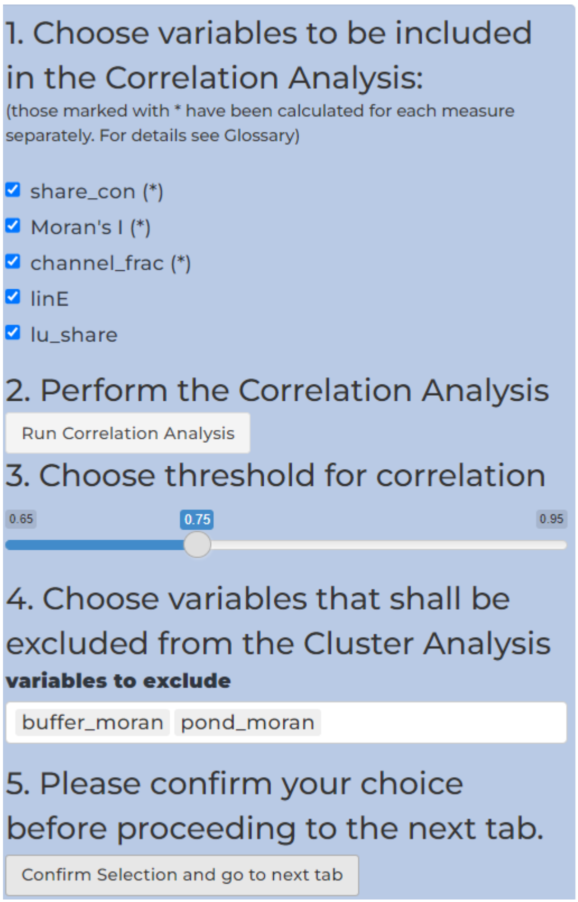

Figure 13: Sidebar of the Correlation Analysis tab

The main panel of this tab is reserved for the correlation figure and a table detailing the degree of correlation among variables. The figure, an example of which is provided in Figure 14, appears after selecting the decision space variables to be considered and clicking on "Run Correlation Analysis". It depicts negative and positive correlation from -1 to +1. In the German CS, the highest correlations can be seen between the share of small measures (hedge, pond, buffer, grass slope) in the catchment area considered for implementation (share_con) and the respective Moran's I. The reason for this lies in the way Moran's I is calculated; its formula includes a denominator that scales the result by the overall variance of the data. When dealing with small land cover types, which typically have fewer data points, the total variance tends to be lower and Moran's is hence more sensitive to small changes in area size and spatial patterns (Hamylton and Barnes, 2018; Huo et al., 2012).

The threshold for correlation that can be set under point 3. is connected to both the table and the selection bar under point 4. The user can examine the variables for which the selected thresholds apply in the table. In the bar under point 4, they can then decide which variables to remove from the subsequent cluster analysis and confirm this by clicking "Confirm Selection and go to next tab". The app then jumps to the next tab.

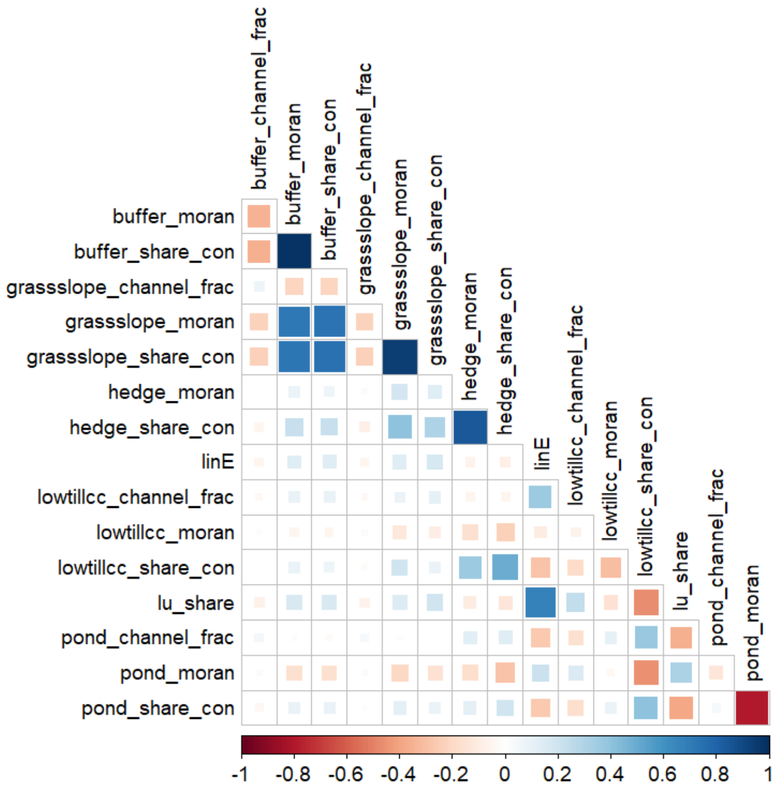

Figure 14: Correlation among decision space variables. Example from CS 1 - Schwarzer Schöps.

4.2.6. PCA & kmeans/kmedoid Tab

In the second tab for manual clustering, the settings for the cluster analysis can be altered and optimised. This tab consists of a sidebar with a list summarising the input variables remaining after the selection in the previous Correlation tab and a summary of the settings selected in the main panel. For easier understanding, the process for selecting cluster settings is numbered, beginning with the plotting aesthetics, followed by the number of principal components to be tested and by the settings for outliers and number of clusters. For the option "Run PCA and Cluster Analysis" to become available on the sidebar, the settings for aesthetics and for the number of principal components have to be confirmed by clicking on the respective buttons. The user can switch between kmedoid and kmeans by clicking the respective checkbox on the sidebar.

All available settings are shown in Table 6. The user can also decide to only test a subset of these settings by selecting "No" in the other setting. To determine the best data ranges for testing and to safe on runtime, it is recommended to first perform a few cluster runs for different settings separately.

The output produced by the clustering is a file called kmeans_data_w_clusters_representative-solutions.csv (or similar depending on which cluster method has been chosen and if outliers have been tested). This file is written to the output folder. If the user wants to reassess the cluster results at a later point and compare them to another run in the app, they will have to save this file in a different location. Otherwise, it might be overwritten by the next cluster run.

The following tab Cluster Analysis always reads from the most recent .csv file in the output folder.

Table 6: Overview of Correlation and Cluster settings that can be optimised in ParetoPick-R

variable description	range		default value
	min	max	
degree of accepted correlation among individual input/decision space variables	0	1	0.7
number of principal components	1	x*	x*
the number of standard deviations a variable within a point must be away from its mean to be considered extreme	0	∞	not tested
the count of variables that exceed the standard deviation threshold within a point	0	∞	not tested
number of clusters	0	∞	15
the accepted maximum share of optima within a cluster that must be extreme for the entire point to be considered an extreme solution	0.1	0.9	not tested

* x is the number of input cluster (decision space) variables remaining after the correlation analysis

4.2.7. Cluster Analysis Tab

This tab shall walk the user through the analysis of the cluster results, help them assess and evaluate cluster performance and understand the representative solutions. The upper part of this tab is divided in two; on the left, a table (see Table 7) details the individual cluster characteristics including cluster size and the objective values of the representative optima.

Table 7: Overview of optima remaining after clustering as provided in tab "Cluster Analysis". Example from CS 1 - Schwarzer Schöps.

cluster number	cluster size	P load at outlet	Ann. min discharge	Crop production	NSWRM costs	opti-mum
0	122	6487	0.1415	59076	49471	615
1	153	5533	0.1452	58934	334854	467
2	130	5342	0.1467	58728	613563	1158
3	210	5910	0.1456	59040	354328	170
4	64	6146	0.1425	58736	32222	951
5	34	6211	0.142	58756	41465	1246
6	118	5804	0.1429	58708	16587	255
7	85	5255	0.146	58707	477086	1449
8	128	6106	0.1437	58961	104414	471
9	160	5890	0.1445	59030	228644	1220
10	110	5269	0.1457	58648	315075	473
11	20	6787	0.1416	59020	6707	200
12	75	6457	0.1416	59022	25668	265
13	29	5249	0.1466	58456	931195	406
14	30	6453	0.1431	59090	103784	1152

On the right, an interactive panel with four checkboxes allows the user to switch between different visualisations. The first plot, selected in the panel shown in Figure 15, shows the Pareto solutions remaining after clustering. Each optimum is representative for one cluster. The user can plot the whole front behind these remaining optima, decide to add the status quo and change the axes, colour and shape like in the previous Pareto front plot.

In the German example, it can be seen that the original full shape of the Pareto front is still visible and well covered by the 15 remaining optima. The optimum representative for the most expensive implementation costs performs well in the environmental objectives but crop production is low. Two optima, representative for the two largest clusters (cluster number 3 and 9 in Table 7) perform very similarly across all objectives. This indicates the need for further analyses and a better separation through an optimisation of the cluster settings (see Appendix Section A1).

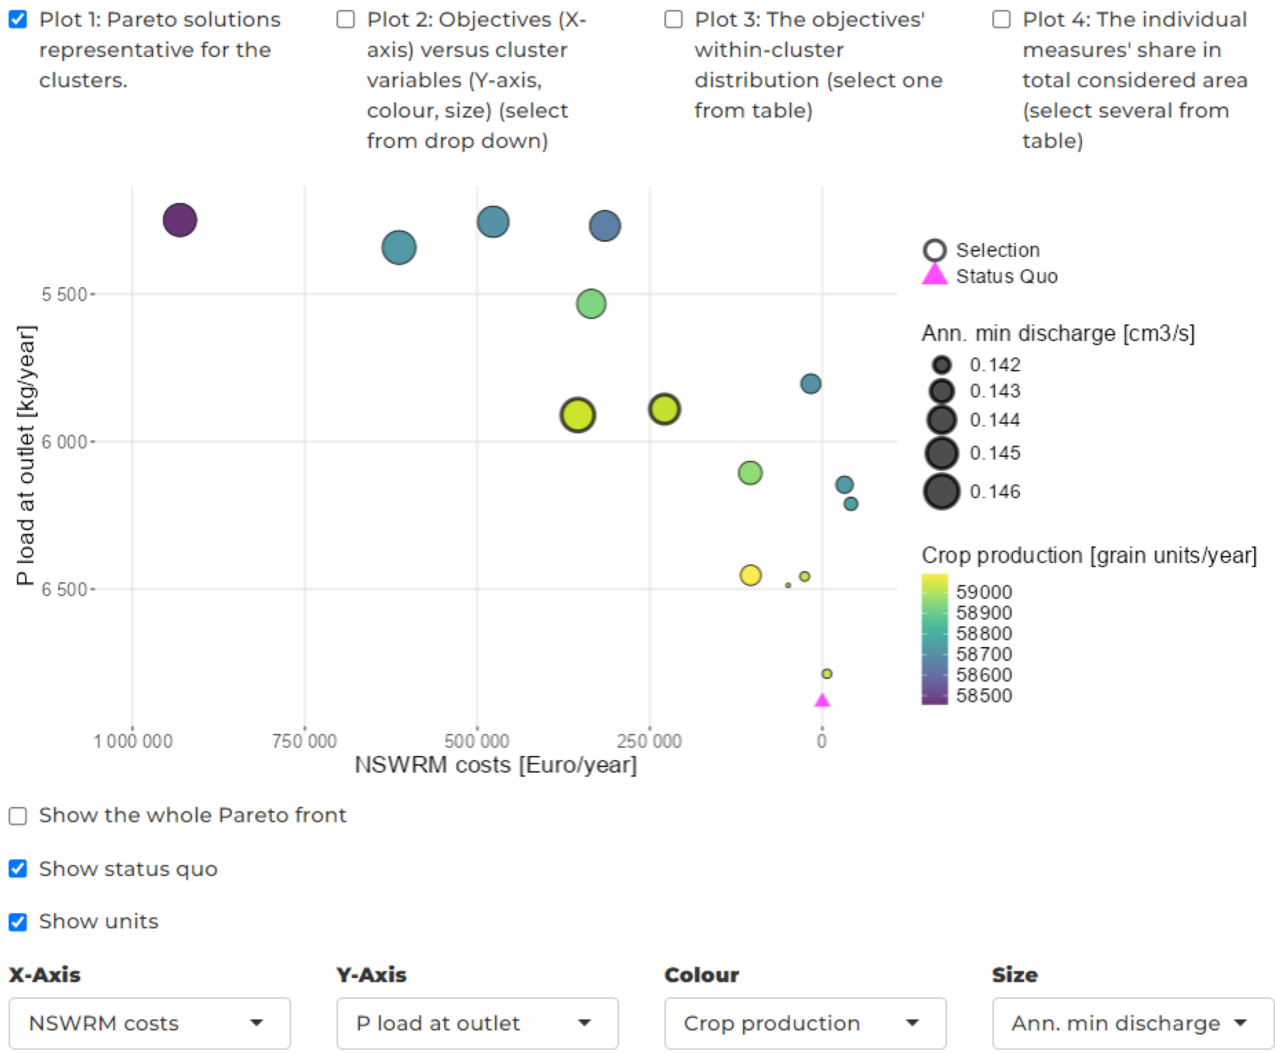

Figure 15: Scatter plot of the Pareto optima remaining after clustering and set of visualisation options. The two largest clusters, 9 (left) and 3 (right) are marked in bold. Example from CS 1 Schwarzer Schöps.

The other visualisations can assist the user to further examine the cluster quality and to gain insight into the procedural reasons for some of the cluster characteristics. The second option is to plot the objectives against the cluster input (decision space) variables. As shown in the example in Figure 16, the former are plotted on the x-axis, while the cluster variables are plotted on the y-axis and as colour and shape. The cluster variables and objectives can be selected from drop down menus. In the example figure, the share of hedges in the area considered for implementation (share_con) is plotted against the implementation costs. Hedges are comparatively expensive which is reflected in the representative optimum with the highest hedge_share_con also displaying the highest implementation costs.

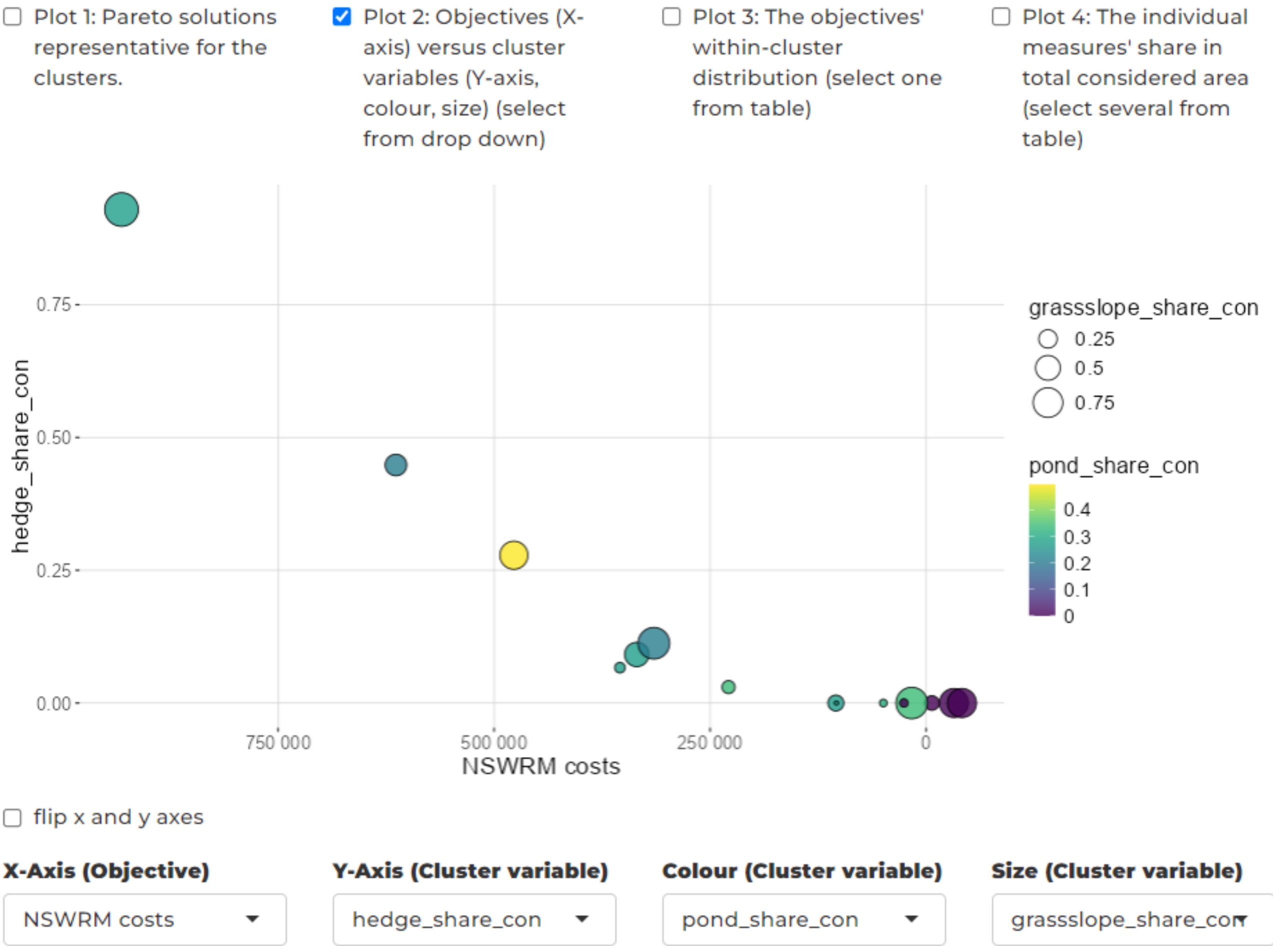

Figure 16: Objective versus cluster variable plot. Example from CS 1 - Schwarzer Schöps.

The third option is to plot the objectives' within-cluster distribution. For easier comparison, the objective ranges are fixed to their maximum range. The cluster can be selected from the table. The example in Figure 17 depicts the largest cluster of the German CS, number 3. The optima representative for this cluster perform well across all objectives, especially in annual minimum flow.

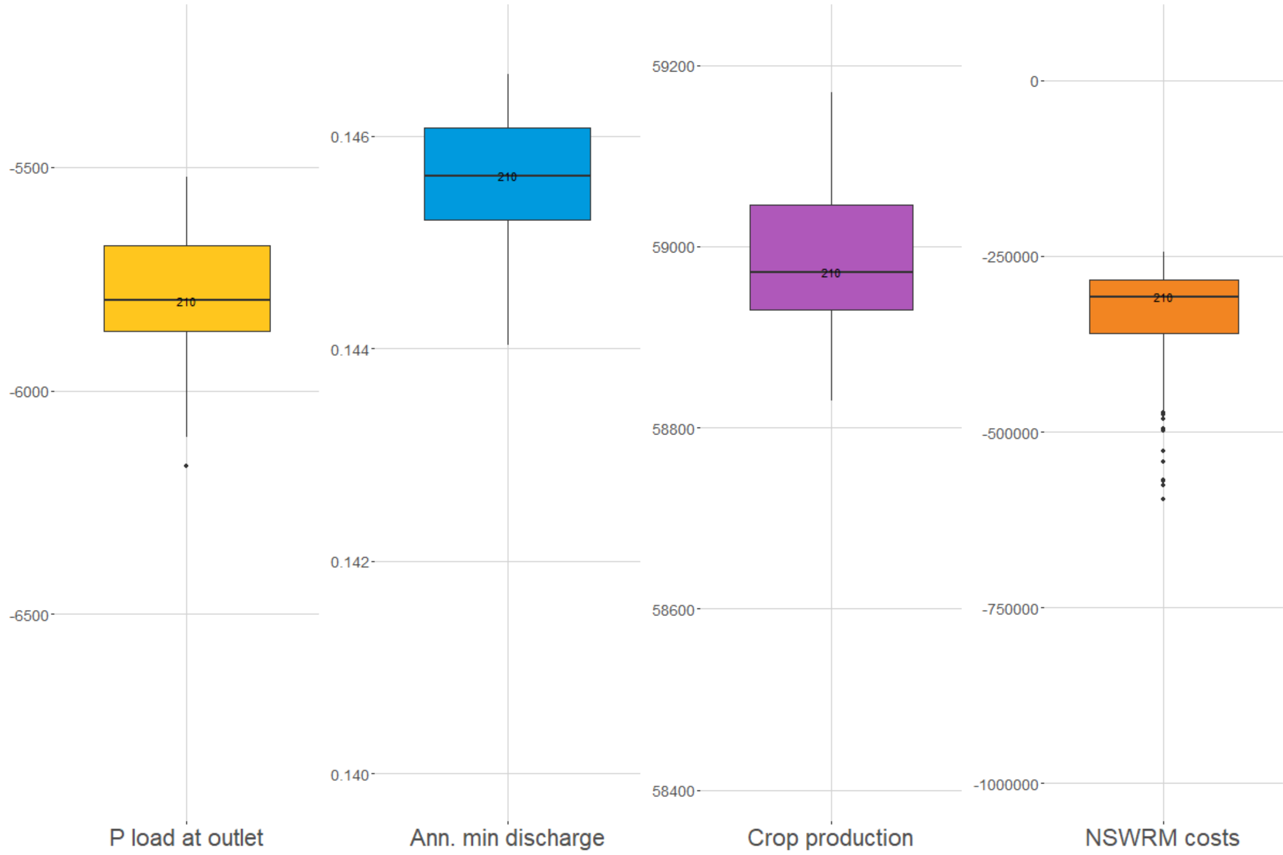

Figure 17: The objectives' within-cluster distribution. Example for the largest cluster number 3 from CS 1 - Schwarzer Schöps.

The fourth option is to compare several clusters selected from the table and to plot the individual measures' share in total considered area. The example in Figure 18 shows the previously discussed clusters 3 and 9 of the German CS. They mainly differ in the share of implemented low till management and the use of hedges. Optima of the smaller cluster 9 implement less of both. Depending on the priorities of the clustering, this suggests the need to adapt it to put less emphasis on these measures as these clusters appear very similar otherwise. In the Appendix A1 (Figure A3), a manual clustering has been performed resulting in a better separation of clusters.

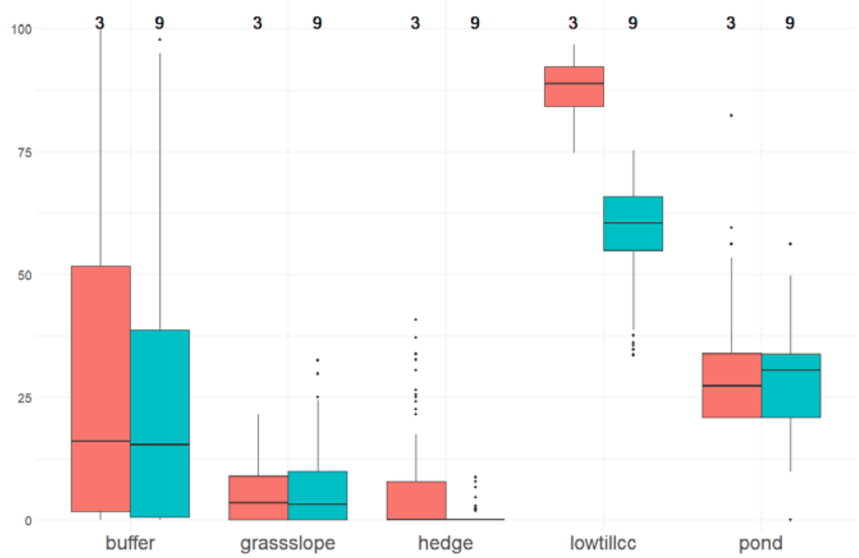

Figure 18: The individual measures' share in total considered area. Example clusters 3 and 9 from CS 1 - Schwarzer Schöps.

Below the table and plot panel, the user can produce maps of representative optima by selecting all clusters/optima of interest in the table and clicking "Plot map of measure implementation plan under selected optima". Figure 19 shows the maps of the two optima representative for the two largest clusters of the German CS. Both of their implementation plans use all measures. Optimum 170 implements low till management throughout the catchment which might account for higher annual minimum flows due to improved infiltration than under optimum 1220 which uses more buffers.

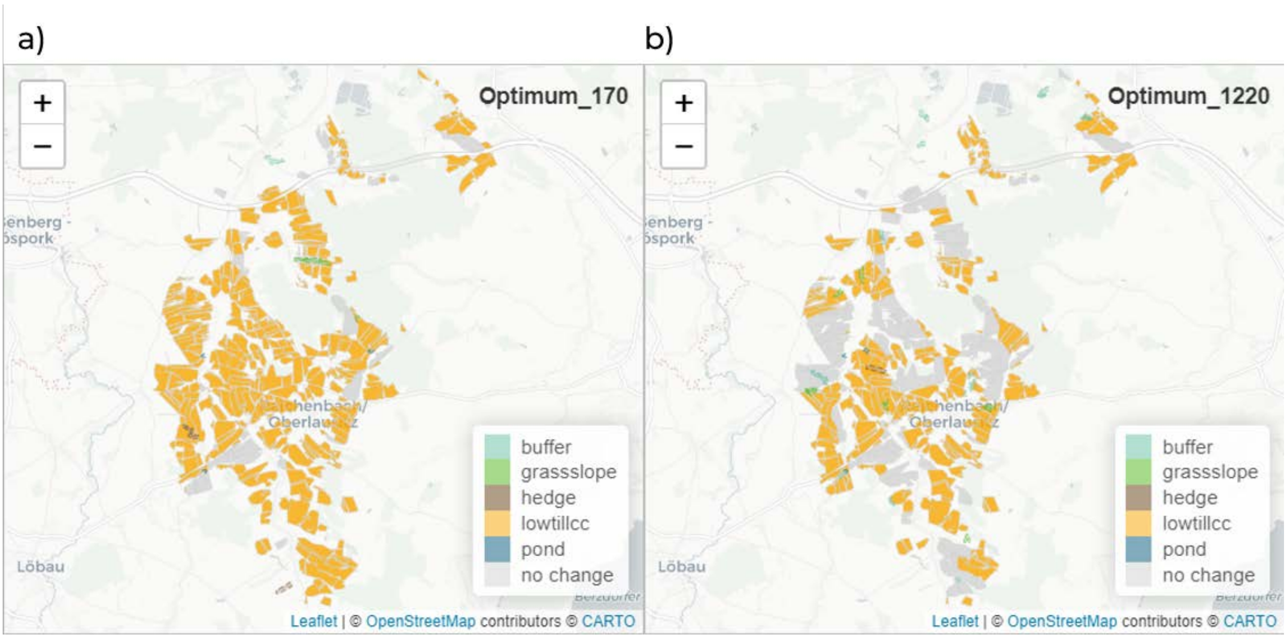

Figure 19: Measure implementation plans of two optima representative for two largest clusters. a) optimum 170 representative for cluster 3 and b) optimum 1220 representative for cluster 9. Example from CS 1 - Schwarzer Schöps.

4.2.8. AHP Tab

This tab allows the user to perform an Analytical Hierarchy Process (AHP). The aim of the AHP is to identify the solution that best meets the preferences of the stakeholders (Section 2.3). This solution is plotted on the Pareto front and the user can produce a map of the respective NSWRM implementation plan. It is possible to select the optimal solution under AHP from the previously determined clusters and thereby combine the two methods.

The tab is structured with a sidebar on the left containing sliders to limit the objective space. The main panel contains cards with pairwise objective comparisons. These can either be opened all at once or one at a time. Like shown in Figure 20, each card provides another scatter plot of the two objectives plotted against each other with the respective R2 and Pearson's r values. Below this plot, the user can assign weights by moving a slider closer to the objective that is more important to them. Above the cards, the currently selected weights are provided in a table. Below the cards, the respective value aligning with the selected weight is given. The user therefore always has an overview of how their decisions affect the overall weighting.

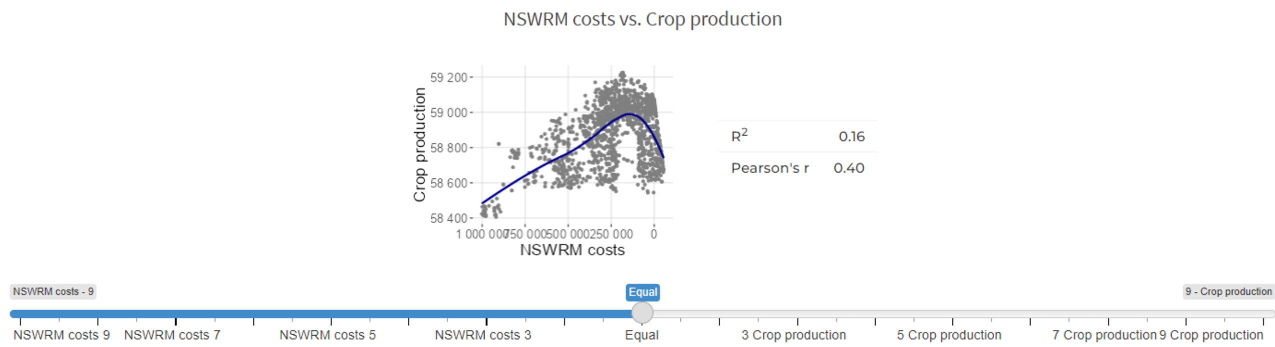

Figure 20: Card for assigning priorities between objectives in pairwise comparison. Example from CS 1 - Schwarzer Schöps.

The general process for performing AHP is to first conduct the pairwise comparisons and assign weights from 2 to 9 to all objective pairs. Choosing the value 9 is consistent with assigning the highest priority to one objective over another. Based on these priorities, AHP calculates weights for individual objectives. Users are notified if their choices are inconsistent (e.g. when objective A > objective B, objective B > objective C, objective C > objective A) and are requested to adapt their selection. The optimal solution under the selected weights is highlighted in the Pareto front and its map can be produced by clicking "Plot map of measure implementation under best option". It is also possible to use the weights to determine the best option among the optima representative for the clusters.

Figure 21 depicts an example of an AHP process. The objective ranges were not restricted. By assigning different priorities among the objectives the weights change, as does the "optimal" solution under AHP. In the example, this is best seen for the NSWRM costs objective. As the weight of this objective decreases (from 1. to 3.), the optimal solution increases in cost and moves to the left along the x-axis. This shift in the optimal solution's position indicates a preference for finding a more economical solution.

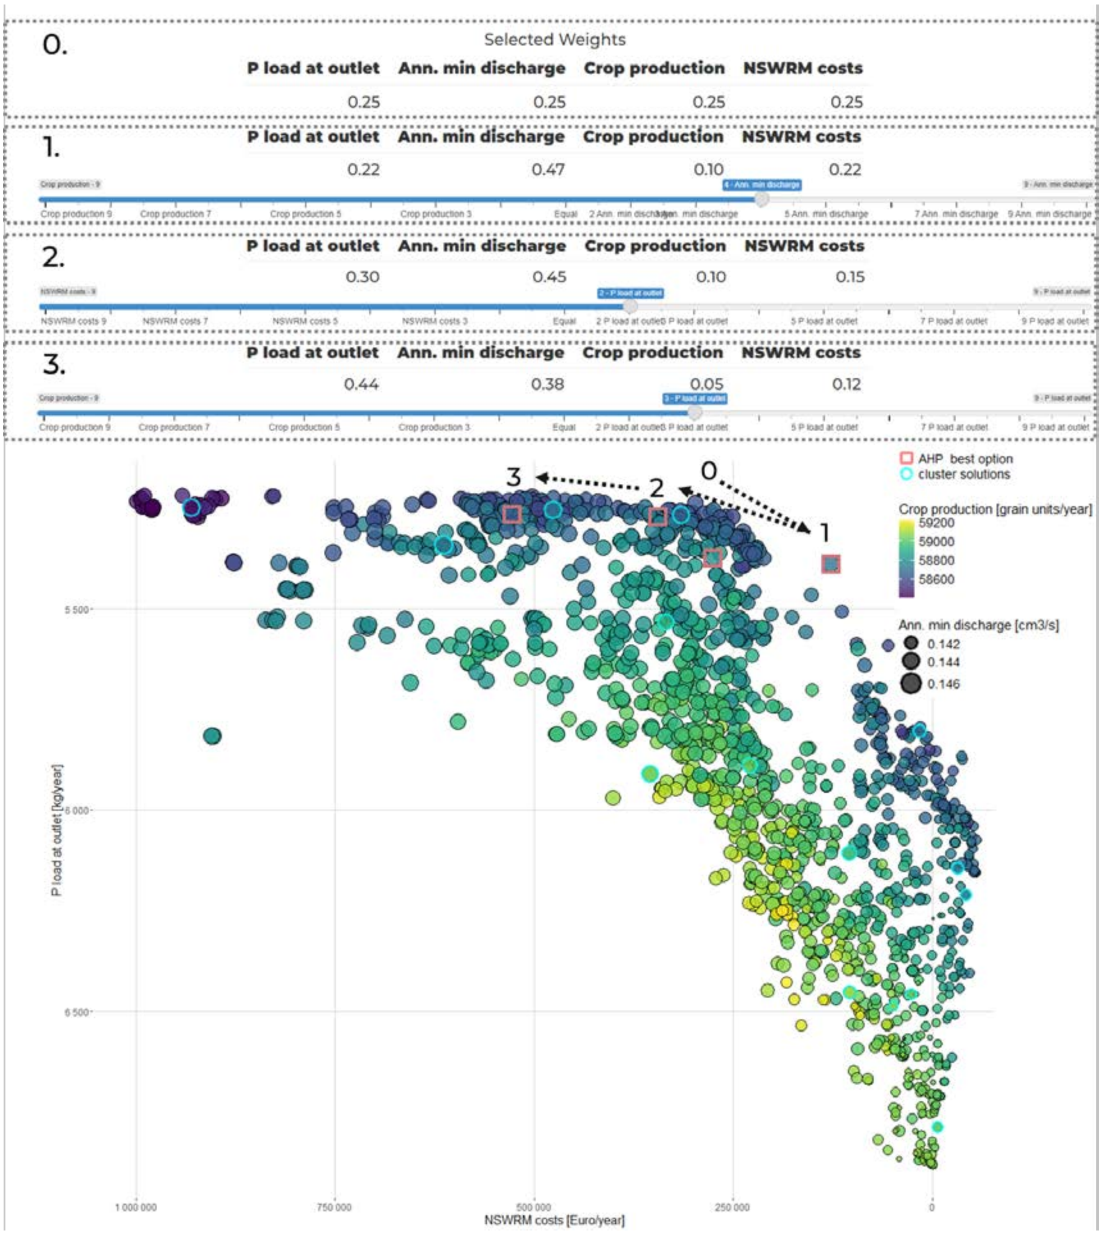

Figure 21: AHP process of assigning preferences and weights. Example from CS 1 - Schwarzer Schöps.

4.3. Export of analyses

The app allows users to save all relevant figures and maps as .png files under a name of their choice. Maps can also be saved as shapefiles. There is furthermore the option to write selected optima to a .csv file so they can easily be reassessed later. In the Visualisation tab, this is the case for the optimum selected in the line plot and in the AHP tab the optimal solution under the selected weights can be written to the .csv.

4.4. Error handling and feedback

While errors and crashes can never be fully prevented, the app was designed with robust error handling and feedback mechanisms to ensure a smooth and user-friendly experience. Key features include input validation and error notification with clear error messages. For example, if the selected slider input is not consistent with the provided data or if the provided data files have the wrong names. Wherever possible, loading spinners have been implemented as progress indicators to inform users about ongoing processes, reducing uncertainty during lengthy computations. The output of the externally called Python scripts is conveyed within the app to assist in error handling and as a fallback output to prevent crashing.

4.5. Additional features

In cases where the app behaves inconsistently and/or crashes despite error handling, a Reset Button on the Data Preparation tab provides the option to delete all previously supplied files as well as the cluster outputs generated within the app. It resets the app to its default "factory" settings and the user can reupload new data in the Data Preparation tab and rerun the data preparation.

4.6. Performance of ParetoPick-R

The tool's performance in terms of speed is largely dependent on the size and complexity of the shapefile. Tests have shown that maps will take longer to render and the visualisation tab in particular will build and load slower if the shapefile contains many small elements, has holes in it and/or has been created by overlaying polygons without proper care for producing clean geometries. To maintain speed and responsiveness, it is recommended to use a streamlined shapefile that contains only the essential polygons and associated information.

5. Application in the case studies

WP5 supported the individual CSs to perform their own MOO through a 3-day workshop in May 2024 and the common optimisation protocol (Strauch and Schürz, 2024). Nevertheless, only six out of 14 CSs completed the optimisation before this report was submitted. The reasons for not completing the optimisation have less to do with the optimisation process itself than with the problems in completing the SWAT+ models. These problems are very diverse, ranging from staff departures to significant data gaps. The current status of all CSs is provided in the Appendix. CSs are encouraged to perform the necessary steps to catch up as soon as possible.

The ParetoPick-R was developed based on intensive feedback from OPTAIN partners. The process of gathering feedback started with a presentation of the tool at the OPTAIN 2024 General Assembly in September in Klaipeda (Lithuania) (World Café format) and continues to date with weekly online video Q&A sessions and the Gitlab and Github repositories, which are kept up-to-date and allow for efficient bug reporting. In fact, many bugs were only discovered by testing with different CS data.

The ParetoPick-R is designed to help CS to manage and understand their optimisation results and to support the dialogue with stakeholders. It is the central tool for the implementation of Task 5.4 “Identification of preferred solutions from actors' perspective”. By May 2025, each CS is asked to conduct 5-10 meetings with stakeholders from different sectors (preferably members of the MARG Core Group, which are already familiar with the project and its modelling approach). Each meeting can be held online and should last no longer than 60 minutes. It should include:

1. an introduction to the optimisation approach,
2. an appropriate visualisation of the Pareto front and its inherent trade-offs and win-win situations,
3. maps and overview tables for example solutions as well as
4. an AHP preference weighting query that should be answered individually by each actor.

More detailed guidelines for these meetings will be prepared by the UFZ in March 2025. The ParetoPick-R will facilitate the whole process and will mainly be operated by a CS lead or modeller. It therefore requires extensive prior testing. Six CS have already done this and were able to provide selected key results in the Appendix.

The UFZ is the main contact for questions related to the optimisation workflow and the ParetoPick-R application. WULS and UFZ are available for questions related to SWAT+ model set-up and UMIL for questions related to the economic model.

6. Outlook

6.1. Integration in the Learning Environment

OPTAINs Learning Environment (LE) is a web resource designed in WP7 to provide in-depth knowledge on NSWRMs based on the project’s main results (Amorsi et al., 2022; Amorsi and Lanceleur, 2023). It serves as a structured guide for researchers and stakeholders seeking to enhance their understanding of NSWRMs, their applications, and the principles behind sustainable water and nutrient management. In OPTAIN, the assessment of the effectiveness of NSWRMs is mainly based on modelling. Therefore, the LE should integrate the modelling results and allow a user-friendly exploration and visualisation of each CS's Pareto-optimal measure implementation plans. As an R-shiny application, the ParetoPick-R can be integrated directly into the LE. However, this requires:

- access to and storage of all necessary CS input data in a ready-to-use format (i.e. for the application within the LE, there should be no need to perform the data preparation step, Section 4.2.2),
- an extensive testing of all ParetoPick-R functionalities relevant to the LE (for each CS),
- Modification in the source code where this is required to provide functionality on the LE website (e.g. related to website interactions).

The functions of ParetoPick-R relevant to the LE and the possibility of direct communication between the two interfaces via various queries and filters (= seamless integration), are still to be defined more precisely in the OPTAIN consortium. Desirable from a WP5 perspective would be a seamless integration of:

- all functionalities of the Visualisation Tab (Section 4.2.3),
- options to filter the solution space based on preferred weights (AHP, Section 4.2.7), and
- the possibility to display and analyse results of at least one clustering (selected by CS leads) in the Cluster Analysis Tab (Section 4.2.7)

To further customise the clustering (Sections 4.2.4-4.2.6), the user is referred to use the full application (i.e. queries within the ParetoPick-R).

6.2. Future development needs & applications beyond OPTAIN

ParetoPick-R's functionalities can be harnessed across a range of applications beyond OPTAIN. The visualisations of the first tab require only one input file structured like pareto_fitness.txt. The app was purposefully designed to allow for such limited data input and (re)creating the file structure of pareto_fitness.txt is straightforward. For all projects where this is an option, especially when only up to four objectives are used, this tab is therefore directly usable, offering versatile capabilities for examining MOO outputs.

For visualisations beyond the Pareto front, more data files are required. While many applications might not be spatial and therefore do not require maps of the decision space, for those were this is relevant, a reproduction of the inputs (e.g. rout_unit.con, hru.con and the HRU shapefiles) required to run ParetoPick-R is near impossible without a SWAT+ model. The cluster process was designed to be performed across variables describing a spatial decision space and are calculated based on SWAT+ and shapefile information. The clustering implemented in ParetoPick-R therefore cannot be performed by other projects working with different data structures.

A functionality currently under development, is the option to not only filter the optima using the objective ranges, but based on measures. The user can then subset the optima according to their preferred share of implemented measures and closer examine the relationship between decision and objective space. Since there is no order of implementation among measures in different locations, the reduction to, for example, 80% of all possible hedges has to be randomised. Different hedges are removed from the set every time the process is repeated. This means, using the sliders for limiting the share of implemented measures results in a similar, but not necessarily in the same, reduced set of optima.

The AHP tab and its intuitive process and visualisations could be valuable for many applications, including those with very different overall objectives than OPTAIN. The interactive sliders for limiting the objective space and for setting priorities assist researchers and other actors in communicating complex and multidimensional decision/objective spaces. Due to the immense relevance of AHP processes across environmental and social applications, a standalone tool based on the AHP tab's source code would be appealing to a wide audience. This application would have to allow for more variable data input formats than applied in OPTAIN and could be extended by a comprehensive set of engaging visualisations.

A few other potential improvements and extensions to ParetoPick-R shall briefly be outlined. For enhancing its applicability and dynamism, ParetoPick-R could be adapted to accept and process a variable number of objectives. Currently, it is required to provide MOO results for four objectives. By using dummy variables, it is possible to process less than four objectives, but allowing a variable number would increase the application's scope and flexibility. More dynamic measure priority allocation would also enhance flexibility. ParetoPick-R hardcodes measures and their respective priority order because particularly the latter is not always known by the users.

Another option would be to implement other visualisations, such as a 3-D plot the user can move and zoom around in. This can be facilitated through a HTML file exported from the app. The R package for producing such file exports, plotly (Sievert, 2020), also provides more options for interacting with visualisations, such as zooms for the Pareto plots. Among other advantages, this would enable an improved distinction and examination of tightly clustered optima. For plotting and processing a high number of optima at high speed, ParetoPick-R relies on ggplot (Wickham, 2016). If these processes should be switched to use plotly instead, the implementation of computational performance improvements such as parallel processing would be required to guarantee an acceptable response time and render speed.

Another interesting extension would be the option to use other methods for clustering that can overcome some of kmeans/kmedoids limitations. One method is fuzzy clustering where individual points (optima) do not have to belong to only one cluster (Miyamoto et al., 2008). Another option would be Density-Based Spatial Clustering of Applications with Noise (DBSCAN) to find clusters of more variable shape, where kmeans assumes a spherical cluster shape (Ester et al., 1996). Like plotly, DBSCAN would necessitate improving computational efficiency, as its limited runtime performance compared to kmeans is well documented (Cebeci and Yildiz, 2015; Panda et al., 2012).

For people with little to no experience in using R, a portable app, for example a simple executable, would be helpful. This can for example be achieved with electron (https://www.electronjs.org/). Such software products are beyond the scope of this OPTAIN deliverable but its Learning Environment will provide sufficiently flexible capabilities.

7. References

Álvarez, M., Moreno, A., Mataix, C., 2013. The analytic hierarchy process to support decision-making processes in infrastructure projects with social impact. Total Quality Management & Business Excellence 24, 596–606. https://doi.org/10.1080/14783363.2012.669561

Amorsi, N., Jacquin, N., Pronost, J., Fribourg-Blanc, B., Semko, J., 2022. Training analysis - Identifying the needs and capacity of relevant target groups for tailoring the OPTAIN Learning Environment to their potential users' requirements. Deliverable D7.5 of the EU Horizon 2020 project OPTAIN. Zenodo.

Amorsi, N., Lanceleur, P., 2023. Learning Environment development strategy (update 1). Deliverable D7.4 of the EU Horizon 2020 project OPTAIN.

Cebeci, Z., Yildiz, F., 2015. Comparison of K-Means and Fuzzy C-Means Algorithms on Different Cluster Structures. Journal of Agricultural Informatics 6. https://doi.org/10.17700/jai.2015.6.3.196

Chang, W., Cheng, J., Allaire, J., Sievert, C., Schloerke, B., Xie, Y., Allen, J., McPherson, J., Dipert, A., Borges, B., 2025. shiny: Web Application Framework for R. R package version 1.8.1.1.

Dolan, J.G., 2008. Shared decision-making–transferring research into practice: the Analytic Hierarchy Process (AHP). Patient education and counseling 73, 418–425.

Ester, M., Kriegel, H.-P., Sander, J., Xu, X., 1996. A density-based algorithm for discovering clusters in large spatial databases with noise, in: Proceedings of the Second International Conference on Knowledge Discovery and Data Mining, KDD'96. AAAI Press, Portland, Oregon, pp. 226–231.

Hamylton, S.M., Barnes, R.S.K., 2018. The effect of sampling effort on spatial autocorrelation in macrobenthic intertidal invertebrates. Hydrobiologia 811, 239–250. https://doi.org/10.1007/s10750-017-3491-x

Huo, X.-N., Li, H., Sun, D.-F., Zhou, L.-D., Li, B.-G., 2012. Combining Geostatistics with Moran's I Analysis for Mapping Soil Heavy Metals in Beijing, China. International Journal of Environmental Research and Public Health 9, 995–1017. https://doi.org/10.3390/ijerph9030995

Krzeminska, D., Monaco, F., 2022. Tailored environmental and socio-economic performance indicators for selected measures. Deliverable D2.2 of the EU Horizon 2020 project OPTAIN. Zenodo.

Miyamoto, S., Ichihashi, H., Honda, K., 2008. Algorithms for Fuzzy Clustering, Studies in Fuzziness and Soft Computing. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-540-78737-2

Monaco, F., Rodriguez, D.G., Szulecka, J., Nesheim, I., Chiaradia, E., 2024. Attractiveness and socio-economic assessment of NSWRMs. Deliverable D4.5 of the EU Horizon 2020 Project OPTAIN.

Moran, P.A.P., 1950. Notes on Continuous Stochastic Phenomena. Biometrika 37, 17–23. https://doi.org/10.2307/2332142

Panda, S., Sahu, S., Jena, P., Chattopadhyay, S., 2012. Comparing Fuzzy-C Means and K-Means Clustering Techniques: A Comprehensive Study, in: Wyld, D.C., Zizka, J., Nagamalai, D. (Eds.), Advances in Computer Science, Engineering & Applications. Springer, Berlin, Heidelberg, pp. 451–460. https://doi.org/10.1007/978-3-642-30157-5_45

Petchrompo, S., Coit, D.W., Brintrup, A., Wannakrairot, A., Parlikad, A.K., 2022. A review of Pareto pruning methods for multi-objective optimization. Computers & Industrial Engineering 167, 108022. https://doi.org/10.1016/j.cie.2022.108022

Piniewski, M., Strauch, M., Plunge, S., Schürz, C., Čerkasova, N., Chiaradia, E.A., Witing, F., 2024. Assessment of NSWRM effectiveness under current and future climate at the catchment scale. Deliverable D4.4 of the EU Horizon 2020 project OPTAIN. Zenodo.

Rousseeuw, P.J., 1987. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. Journal of Computational and Applied Mathematics 20, 53–65. https://doi.org/10.1016/0377-0427(87)90125-7

Saaty, T.L., 1977. A scaling method for priorities in hierarchical structures. Journal of Mathematical Psychology 15, 234–281. https://doi.org/10.1016/0022-2496(77)90033-5

Schürz, C., Čerkasova, N., Farkas, C., Nemes, A., Plunge, S., Strauch, M., Szabó, B., Piniewski, M., Weiland, S., Shore, M., Cüceloglu, G., Czelnai, L., Idzelytė, R., 2022. Modelling protocols. Deliverable D4.2 of the EU Horizon 2020 project OPTAIN. Zenodo. https://doi.org/10.5281/zenodo.13981318

Sievert, C., 2020. Interactive web-based data visualization with R, plotly, and shiny. Chapman and Hall/CRC.

Strauch, M., Schürz, C., 2024. Common optimisation protocol. Deliverable D5.1 of the EU Horizon 2020 project OPTAIN. Zenodo. https://doi.org/10.5281/zenodo.11473792

Suarna, N., Wijaya, Y.A., Mulyawan, Hartati, T., Suprapti, T., 2021. Comparison K-Medoids Algorithm and K-Means Algorithm for Clustering Fish Cooking Menu from Fish Dataset. IOP Conf. Ser.: Mater. Sci. Eng. 1088, 012034. https://doi.org/10.1088/1757-899X/1088/1/012034

Thungngern, J., Sriburi, T., Wijitkosum, S., 2017. Analytic hierarchy process for stakeholder participation in integrated water resources management. Engineering Journal 21, 87–103.

White, S.E., Witing, F., Wittekind, C., Volk, M., Strauch, M., 2025a. PyretoClustR (executable version). https://doi.org/10.5281/zenodo.14761002

White, S.E., Witing, F., Wittekind, C.I.H., Volk, M., Strauch, M., 2025b. Distilling the Pareto Optimal Front into Actionable Insights. https://doi.org/10.2139/ssrn.5135378

Wickham, H., 2016. ggplot2, Use R! Springer International Publishing, Cham. https://doi.org/10.1007/978-3-319-24277-4

Appendix - application in the OPTAIN case studies

A1 - Case Study 1 - Schwarzer Schöps, Germany

Full Pareto front

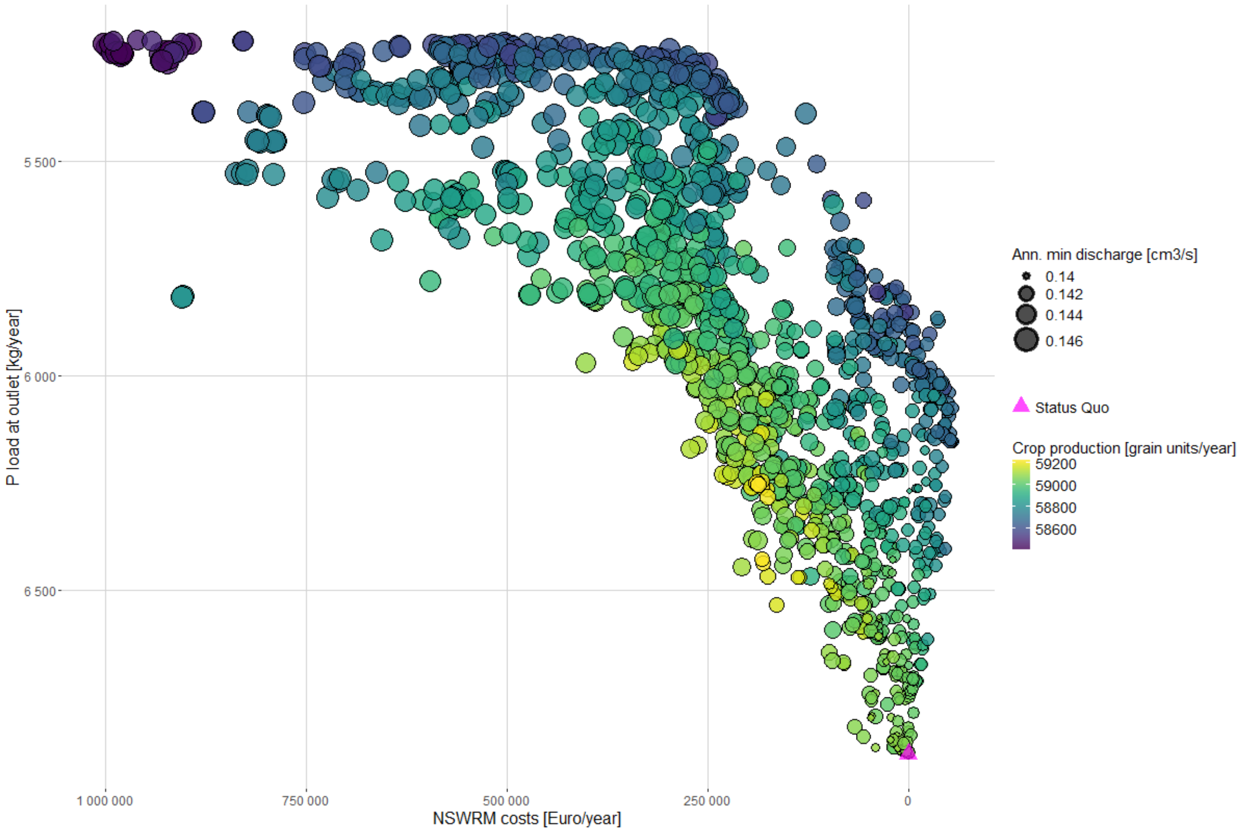

Figure A1: Scatter plot of the full Pareto front. Example from case study 1 Schwarzer Schöps.

Plotting the two objectives with the highest absolute ranges as x and y axes shows the trade-off between the two; in optima with lower P loads, implementation costs tend to be higher. However, this trade-off has a strongly convex shape, indicating a high potential for 'good' compromise solutions. For example, it may be possible to reduce average annual P loads to less than 6000 kg with very low NSWRM costs or even savings (at catchment scale). These solutions (P loads < 6000 kg/year and costs around 0) in turn have a trade-off with the other two objectives, minimum annual runoff and crop production. However, the overall magnitude of change and therefore the risk of loss for these two objectives is comparatively small (<5%).

Frequency Analysis

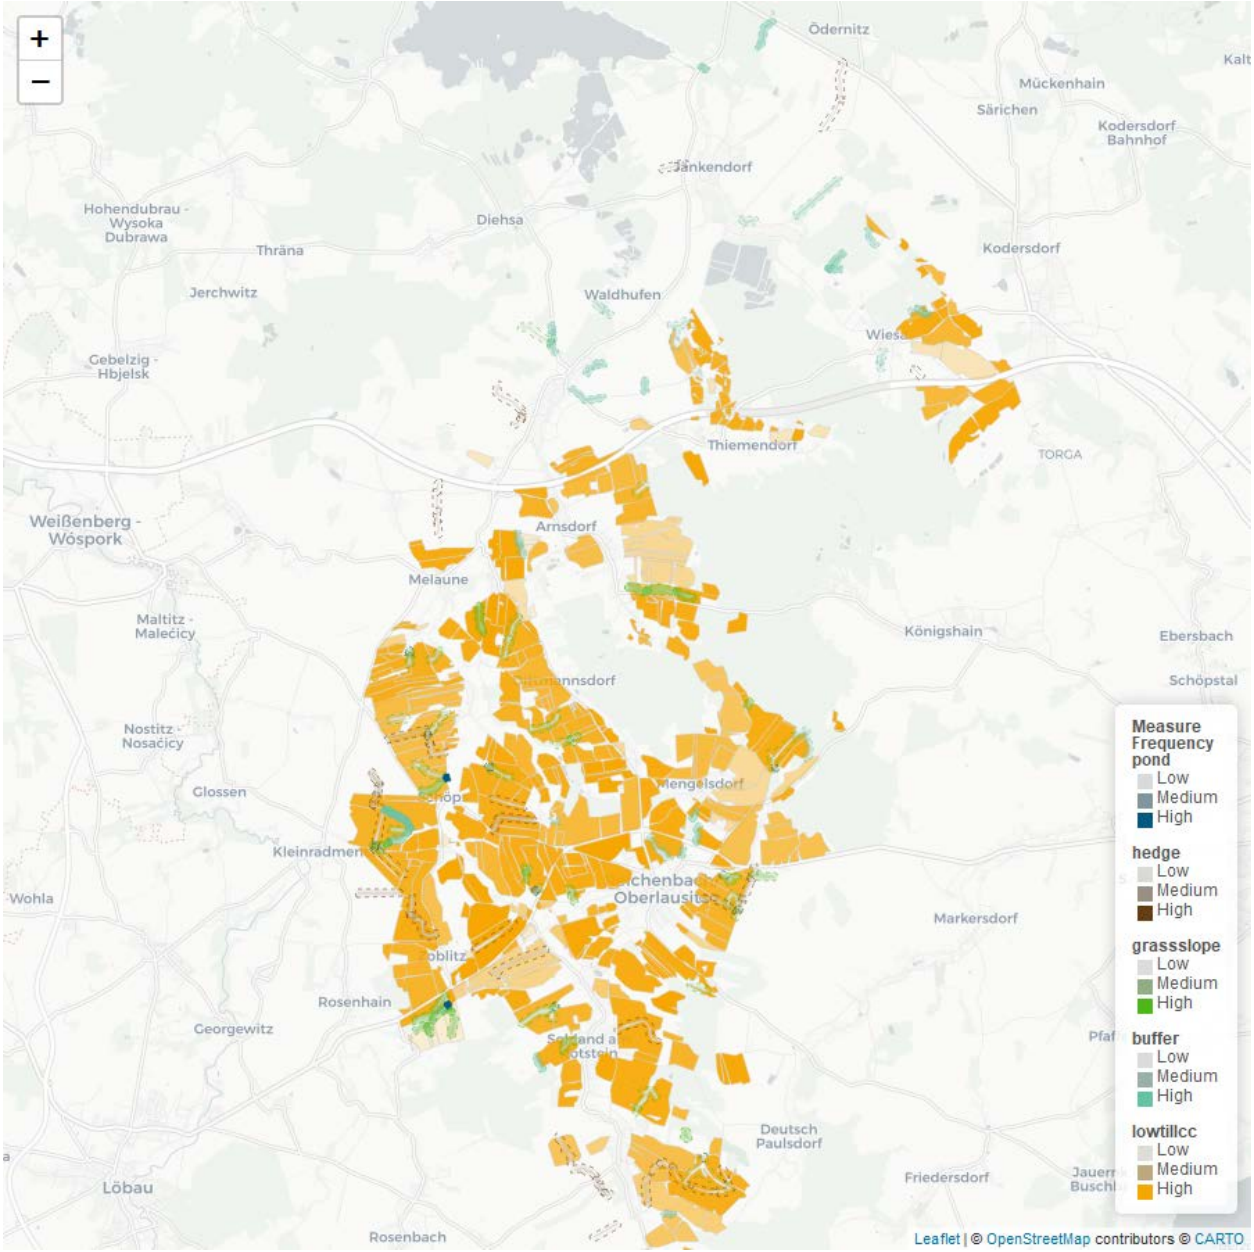

Figure A2: Summary of all measure implementation plans. Frequency with which individual HRUs are activated in those optima performing beyond 0.5 across all objectives. Example from case study 1 Schwarzer Schöps.

Figure A2 shows an example frequency plot for all optima with above-average (>0.5 scaled) performance across all objectives. Hedges are not frequently implemented, likely because they are expensive. One buffer in the west of the catchment appears to be important for overall good performance, as are two very frequently used ponds. Lowtill is frequently implemented across these optima and some areas, such as the fields in the west of the catchment, appear more important for achieving these objective ranges than others.

Cluster Results

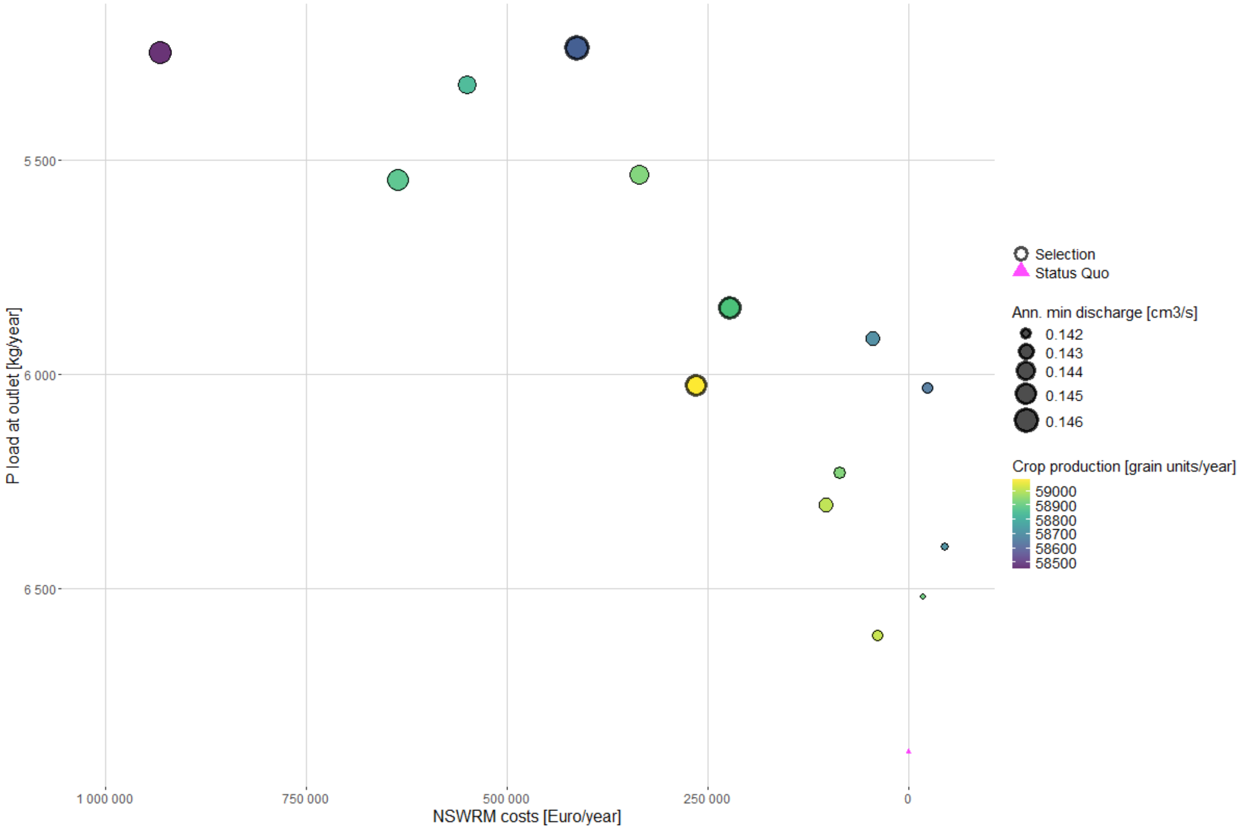

Figure A3: Scatter plot of the Pareto optima remaining after clustering. The three largest clusters (0, 2, 5) are marked in bold (from left to right). Example from case study 1 Schwarzer Schöps.

The clustering has been performed manually with kmeans, testing only the number of clusters (13-18). The result is 14 well separated clusters. The original full shape of the Pareto front is still visible and well covered by the remaining representative optima. The overall relationships among the objectives are also still distinguishable. The optimum representative for the cluster with the highest implementation costs (cluster number 0) performs well in the two environmental objectives but implies below average crop production.

Cluster Analysis

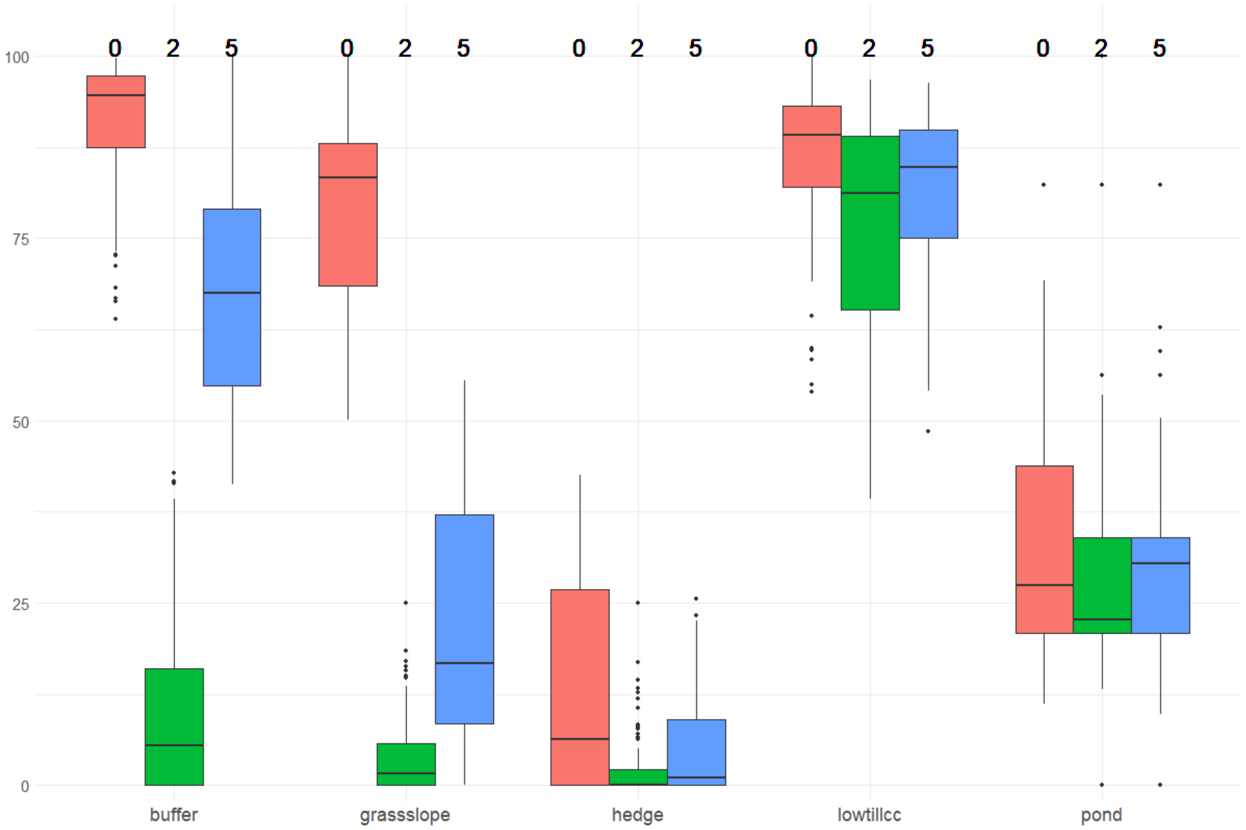

Figure A4: The individual measures' share in total considered area. Example clusters 0, 2 and 5 from case study 1 - Schwarzer Schöps.

The strongest differences between the three largest clusters lie in the share in implemented buffers and grass slopes. Optima in cluster 2 use the lowest number of buffers, grass slopes and hedges. Accordingly, cluster 2 contains optima with the highest crop production and worst performance in the two environmental objectives. Across the three clusters, optima in cluster 0 implement most measures. Accordingly, the representative optima for cluster 0 is the most expensive and displays the best performance in the two environmental objectives. Similar shares in ponds are used across these three clusters.

A2 - Case Study 2 - Petite Glâne, Switzerland

The optimisation study for CS 2, Petite Glâne, Switzerland, was conducted using four performance indicators: P load the catchment outlet, low flow days the catchment outlet, average soil moisture at 30cm and crop yield, represented as grain units.

Full Pareto front

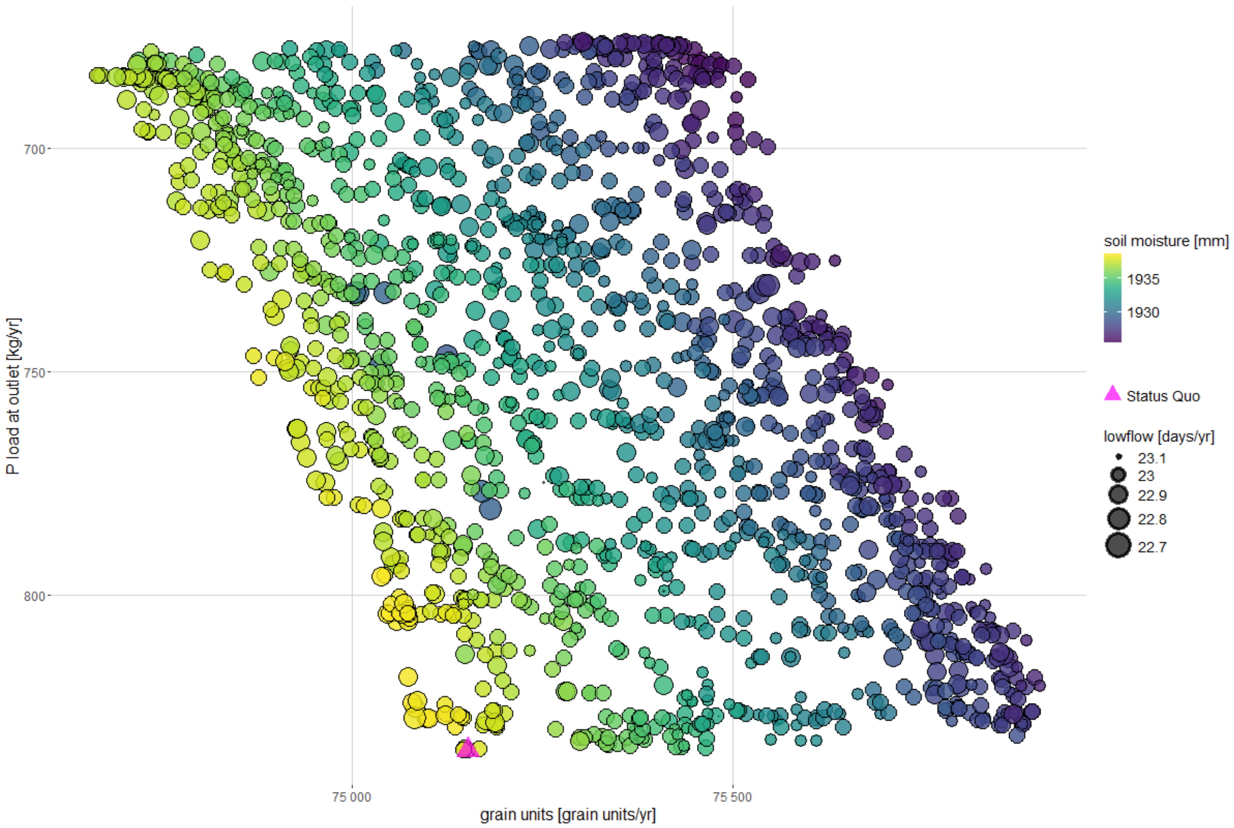

Figure A5: Scatter plot of the full Pareto front. Example from case study 2 Petite Glâne, Switzerland.

The full Pareto front for CS 2, Petite Glâne, (see Figure A5) illustrates the trade-off between P load at the outlet (x-axis) and grain units (y-axis), where lower P loads correspond to fewer grain units. Another trade-off is observed between grain units and soil moisture (colour); with increasing grain units, soil moisture decreases.

Frequency Analysis

Figure A6 presents an example frequency plot for all optimal solutions with above average (>0.5 scaled) performance across all objectives. It can be observed that most measures are implemented in the upper or lower zone of the catchment, while the middle zone remains largely unchanged, resembling the status quo.

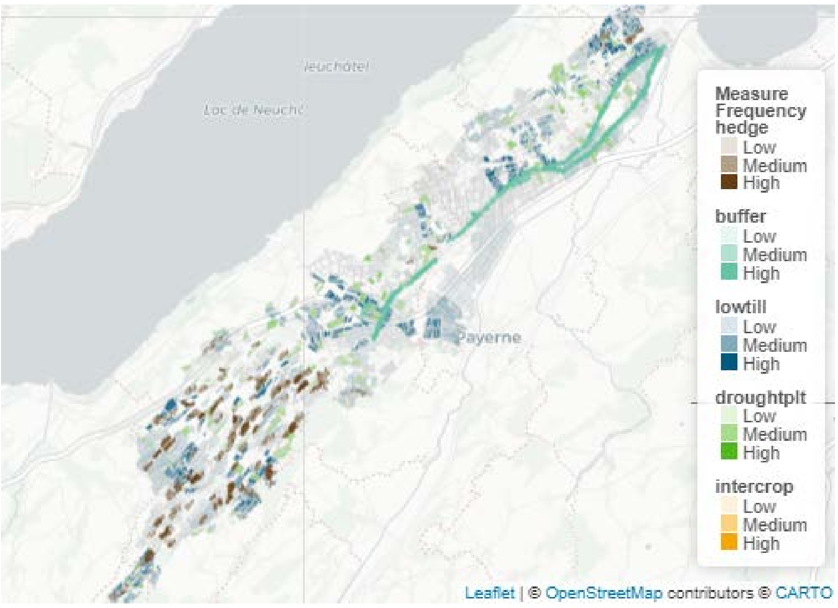

Figure A6: Summary of all measure implementation plans. Frequency with which individual HRUs are activated in those optima performing beyond 0.5 across all objectives. Example from case study 2 Petite Glâne.

Cluster Results

Figure A7: Scatter plot of the Pareto optima remaining after clustering. The three largest clusters (0, 5, 7) are marked in bold. Example from case study 2 Petite Glâne.

The clustering, performed with default settings, produced 15 well-separated clusters. The original Pareto front shape remains visible and well represented by the remaining optima. The three largest clusters are nr. 0, 5 and 7. Cluster 0 achieves high grain unit performance, however, it also results in high P loads at the outlet despite slightly improving upon the status quo. Clusters 5 and 7 reduce P loads while also slightly increasing grain units compared to the status quo.

Cluster Analysis

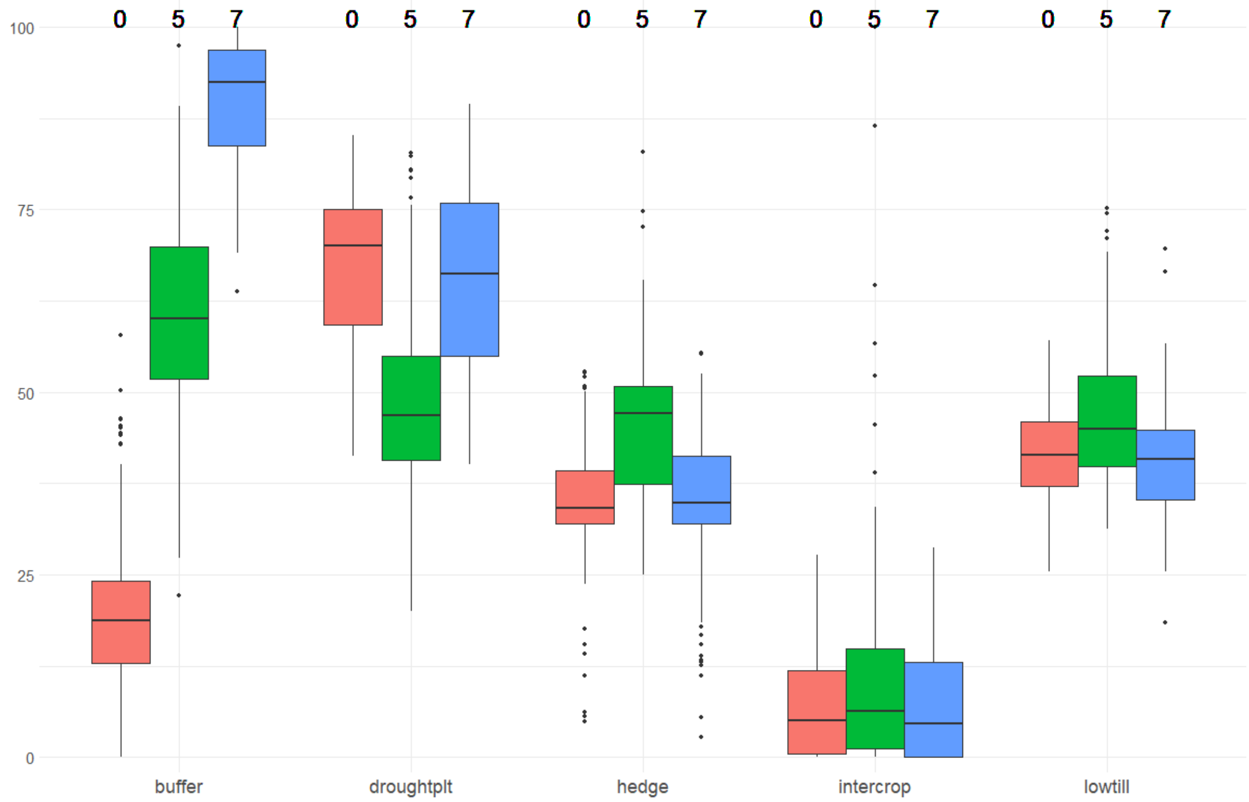

Figure A8: The individual measures' share in total considered area. Example clusters from case study 2 Petite Glâne.

The biggest difference among clusters 0, 5 and 7 is the share of implemented buffers. While the optima in cluster 0 include almost no buffers, cluster 7 utilises nearly the entire considered area for buffers. Figure A8 also shows that in all three clusters, the share of intercrop remains very low. This is also reflected in the frequency map (Figure A6), where no intercrop measure is not visible.

Outlook

In the future, we plan to replace grain unis with gross margin; however, the gross margin calculation is not yet finalised. Moreover, we might replace some other indicators; since for example low flow days displays very little variation, it could be replaced by the mean flow at the outlet, while P loads could be replaced by soil erosion, another indicator selected by the stakeholder in the 2nd MARG Workshop.

A3 - Case Study 3 - Felső-Válicka, Hungary

To determine the optimal spatial distribution of the NSWRMs in CS 3, the following environmental and economic indicators in the optimisation calculations were used: nitrogen (N) load in the channel (kg/a), average soil moisture content in the top 30 cm of soil at the agricultural parcels within the catchment (mm), grain yield from the agricultural parcels of the catchment (unit/a), and the implementation and maintenance costs of the NSWRMs (Euro/a).

Full Pareto front

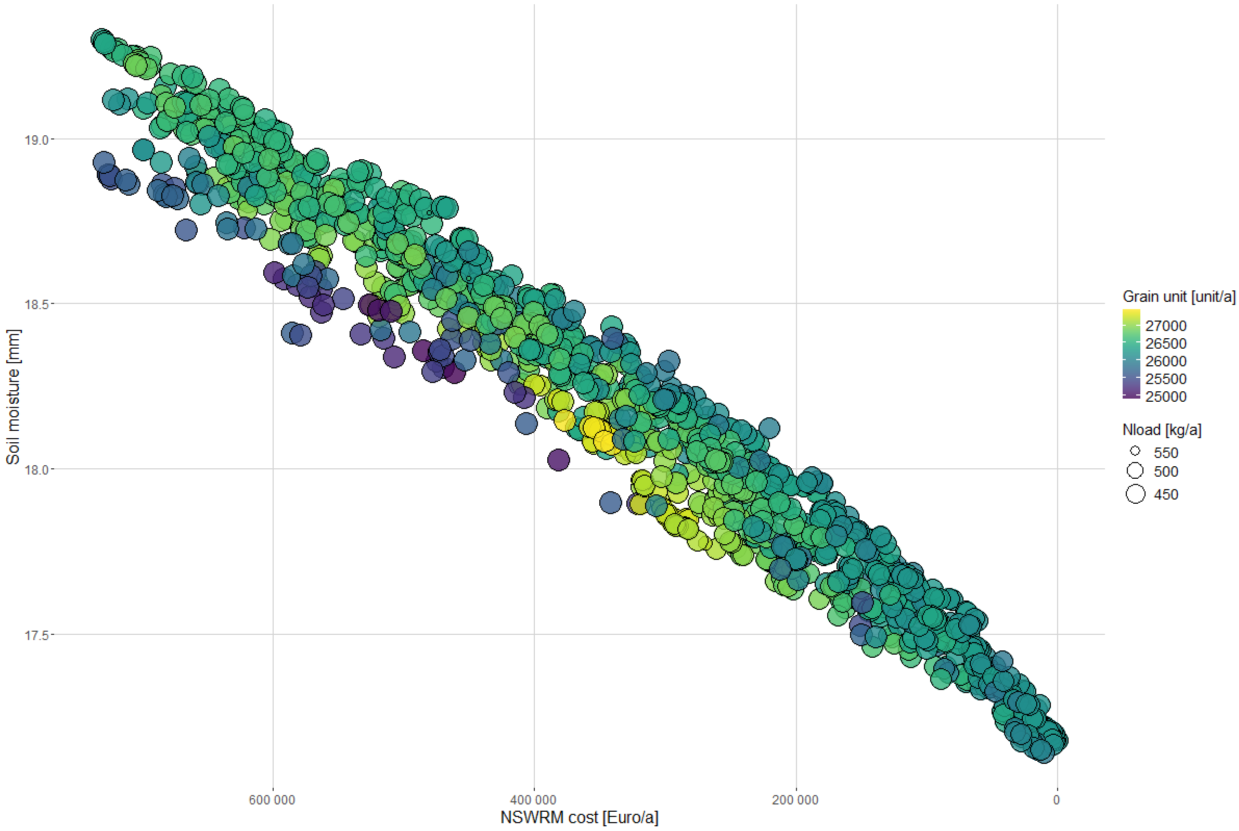

Figure A9: Scatter plot of the full Pareto front of the case study 3 Felső-Válicka.

The full pareto front (Figure A9) shows an inverse correlation between soil moisture content in the top 30 cm and NSWRM costs. More efficient reduction of N load also increases NSWRM costs. A higher grain unit does not exclude the possibility of reducing N load but requires the implementation of measures with higher costs (e.g., Figure A12: representative solution 5).

Frequency Analysis

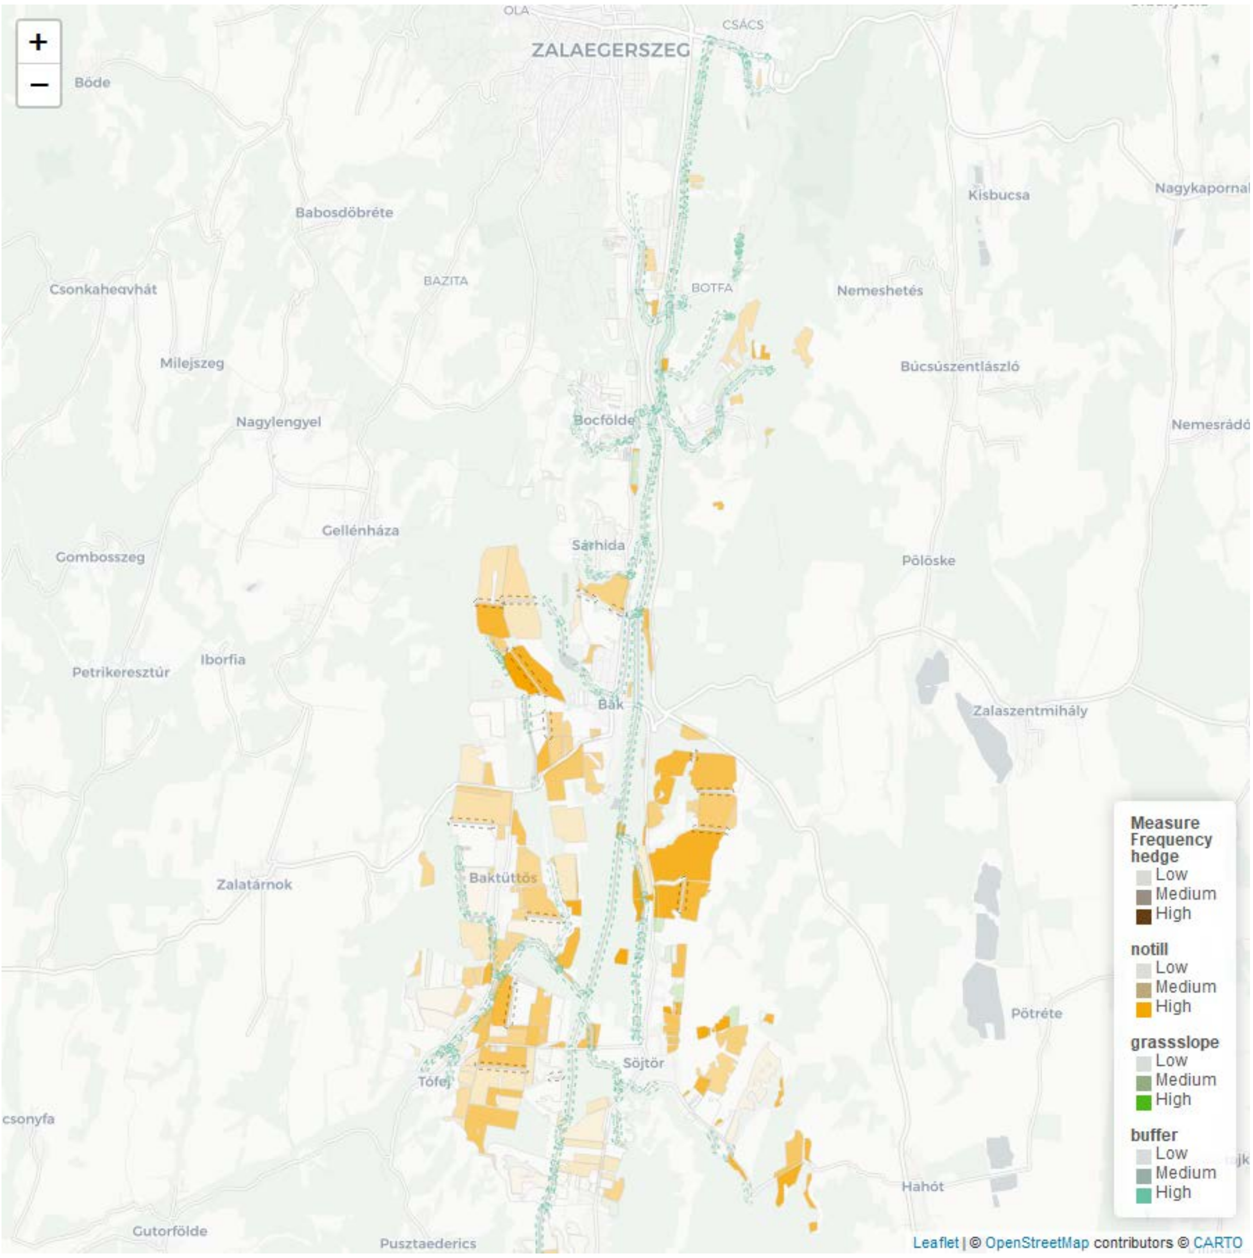

Figure A10: Summary of all measure implementation plans. Frequency with which individual HRUs are activated in those optima performing beyond 0.5 across all objectives. Example from case study 3 Felső-Válicka.

The overall frequency of measures ranges from medium to low for field-dividing hedges, the conversion of arable land to grassland (grass slope), and riparian buffers. In contrast, no-till management with cover crops exhibits not only low and medium but also high frequencies (Figure A10).

Cluster Results

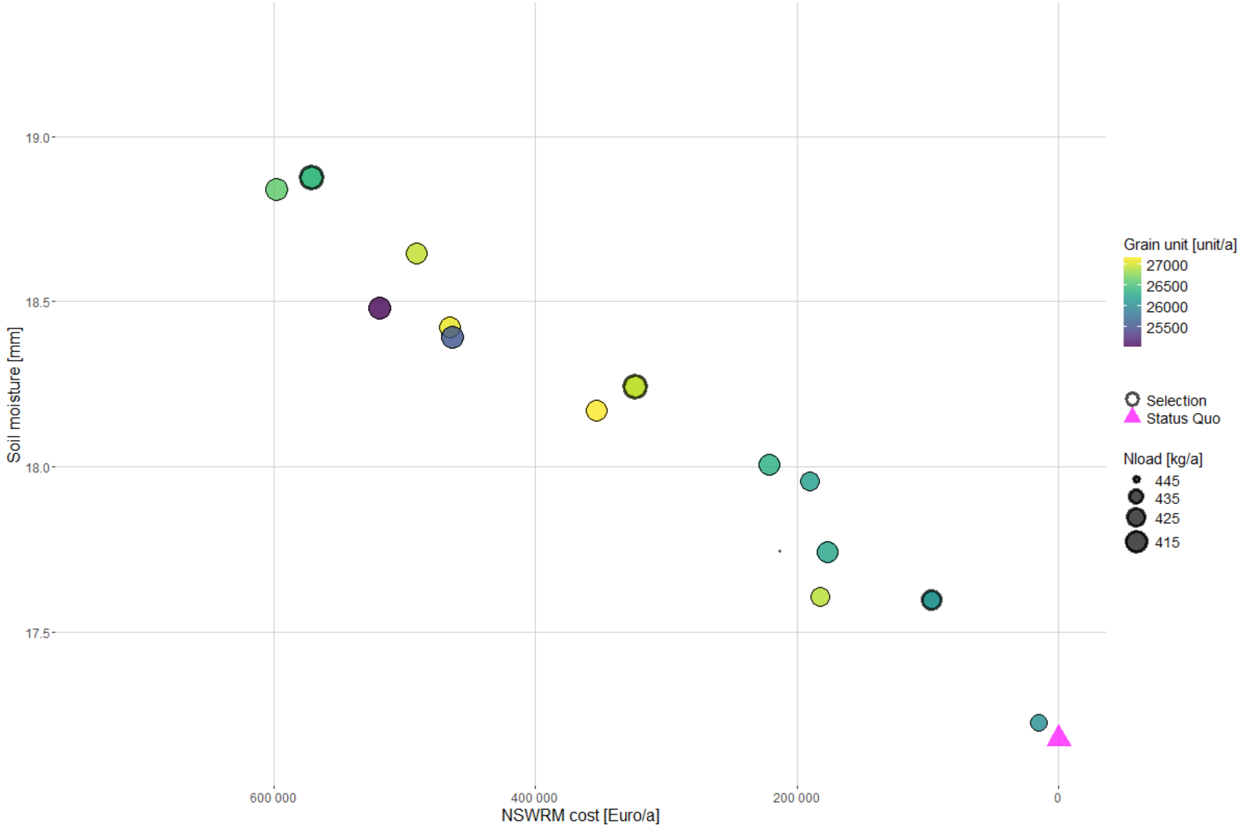

Figure A11: Scatter plot of the Pareto optima remaining after clustering. The three largest clusters (3, 6, 0) are marked in bold (from left to right). Example from case study 3 Felső-Válicka.

Using the default settings for the k-means clustering step, 15 clusters were specified. Figure A11 presents the optimal solutions closest to each cluster's centroid. Cluster 3 exhibits the highest soil moisture content in the top 30 cm, lower N load, a high grain unit, and moderate NSWRM costs (around 600,000 Euro/year per catchment). In Cluster 6, the grain unit remains high, while soil moisture content, N load, and NSWRM costs are at moderate levels. Cluster 0 is characterised by a moderate grain unit and N load, low NSWRM costs, and lower soil moisture content. Clusters 3 and 6 demonstrate greater efficiency in reducing N load compared to Cluster 0.

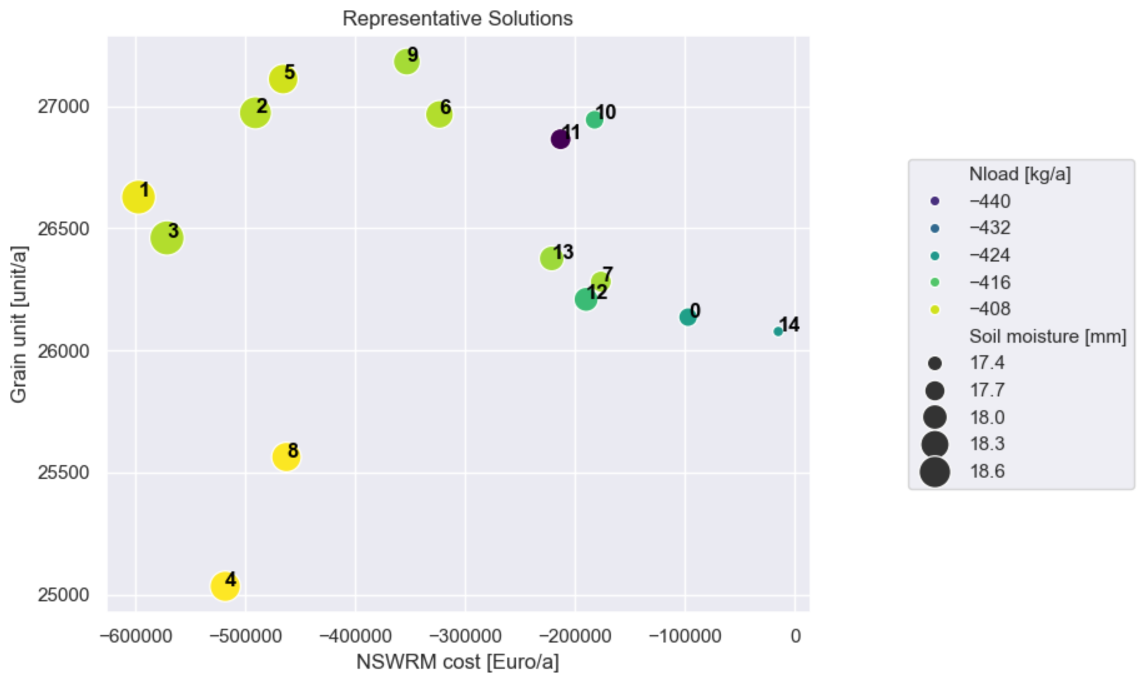

Figure A12: Relationships between the environmental and economic indicators in the case of the representative solutions of the differentiated 15 clusters.

<u>Cluster Analysis</u>

The boxplot diagram of the three largest clusters (0, 3 and 6) highlights the varying proportions of NSWRM application (Figure A13). In the case of riparian buffers and field-dividing hedges implementation, as well as arable land conversion to grassland (grass slope) there are no significant differences among the three selected clusters. Only a small area is allocated to riparian buffers, while field-dividing hedges are not included in any of the optimal solutions. Approximately 6–7% of the arable land is optimally converted into grassland.

The key distinction between the clusters lies in the extent of no-till with cover crops management application. Cluster 3 has the highest proportion of no-till with cover crops management, covering around 70-80% of the arable land. This measure was found to be the most efficient in increasing the soil moisture content, explaining the high moisture levels observed in Cluster 3. In contrast, Cluster 6 has a no-till area share of 37–50%, while Cluster 0 has only around 10–20%.

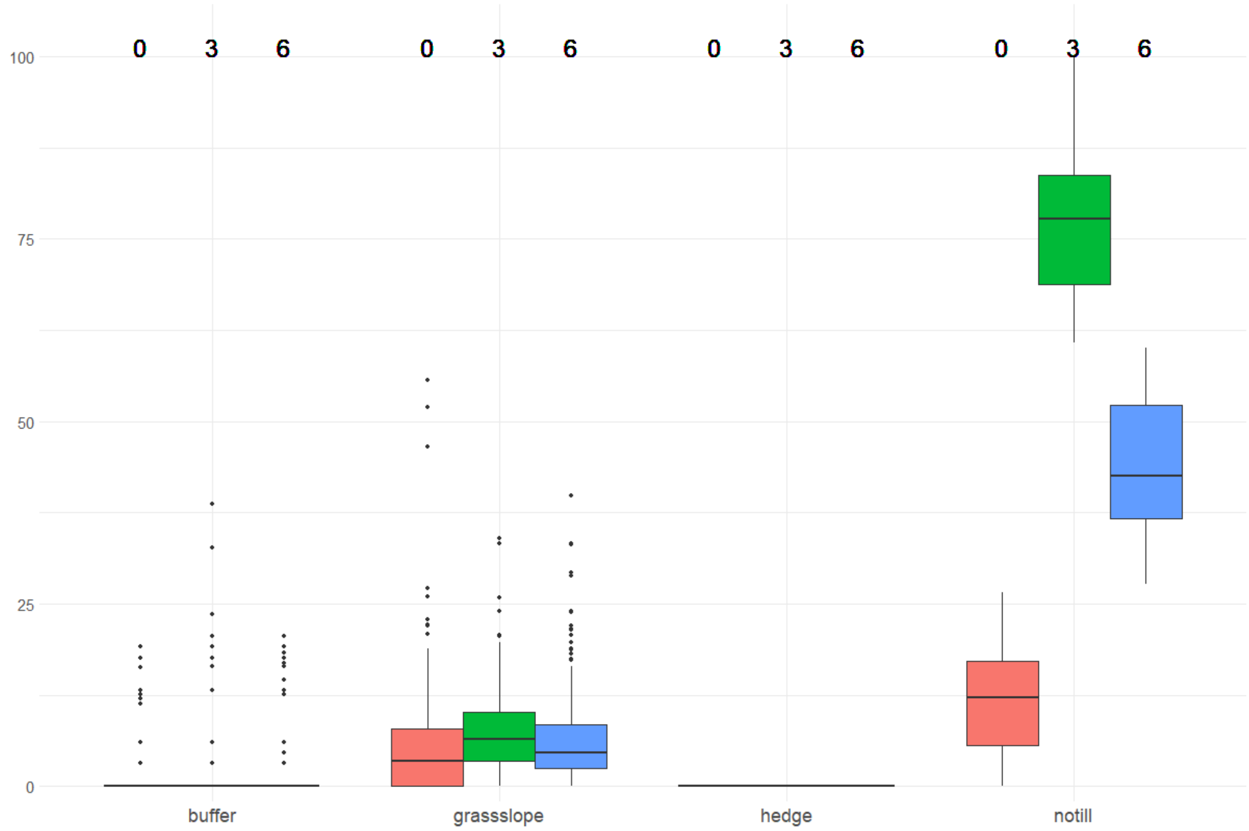

Figure A13: The individual measures' share in total considered area. Example clusters 0, 3 and 6 from case study 3 Felső-Válicka.

<u>Outlook</u>

For the presentation of the optimisation results at the MARG interviews, NSWRM costs shall be replaced with gross margin, as this economic indicator might be more informative for farmers and farm advisors.

A4 - Case Study 4 - Upper Zgłowiączka, Poland

The SWAT+ model and NSWRM scenarios are ready for the optimisation step. However, testing the NSWRMs scenarios had shown that they had no or adverse effects. It was hence decided to update the measures and their representation in the models and run the scenarios again. The CS team is currently setting up the implementation to run CoMOLA.

A5 - Case Study 5 - Pesnica, Slovenia

The SWAT+ model setup for CS5 was finished on 12 February 2025 (Projcet partner WULS). The team also analysed the effectiveness of NSWRMs using the calibrated SWAT+ model and quantified it for four scenarios, which were simulated with SWAT+ using the SWATmeasR package: riparian buffers, crop rotation, cover crops, and ponds.

The performance of the tasks (running the optimisation using CoMOLA, then applying the post-processing using ParetoPick-R) is behind schedule due to changes in the partner's team (UL). The leading modeller left the research group. We are starting the process of using CoMOLA and ParetoPick-R now (month 54) and plan to finish the work by the end of March 2025.

A6 - Case Study 6 - Kebele Kobiljski, Slovenia

The SWAT+ model setup and analysis of NSWRM effectiveness using the SWATmeasR package for CS6 will be developed by 31 March 2025 with the help of the project partner, WULS. The UL team will develop the optimisation using CoMOLA and apply the post-processing using ParetoPick-R by the end of April 2025.

The tasks' performance is behind schedule due to changes in the partner's team (UL). The leading modeller left the research group. We plan to finish the work by the end of April 2025.

A7 - Case Study 7 - La Wimbe, Belgium

The optimisation study for CS 7, Wimbe, Belgium, was conducted using four performance indicators: total sediment load [tons/yr], frequency of low flow days at the catchment outlet, average maximum daily discharge of each year [m^3/s] at the catchment outlet, and crop yield, represented as grain units. The measures implemented in the model were riparian buffers, riparian forest buffers, afforestation, field hedges, wetlands, floodplain restoration.

Full Pareto front

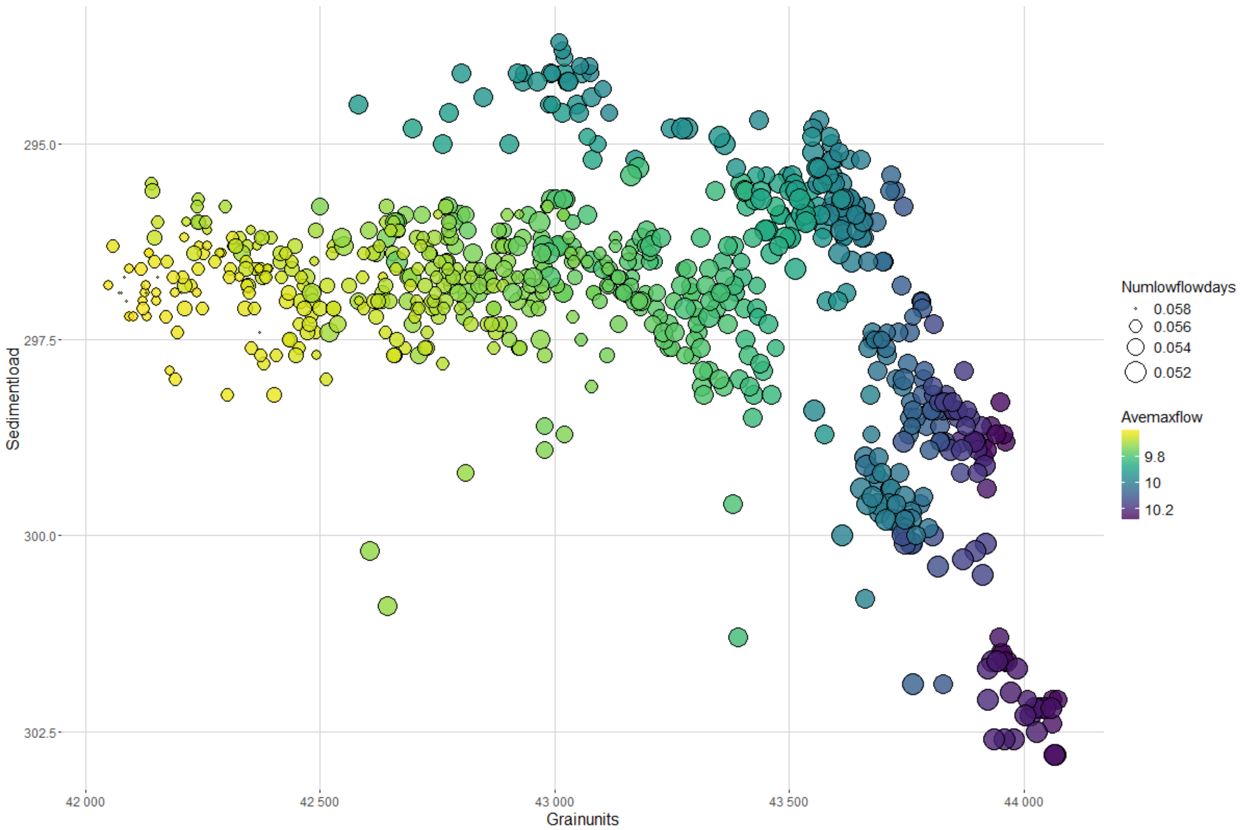

Figure A14: Scatter plot of the full Pareto front. Example from case study 7 Wimbe.

The full Pareto front for CS 7, Wimbe, (see Figure A14) illustrates a trade off between sediment load (y-axis) and grain units (x-axis), wherein, as the measures are implemented, lower sediment load corresponds to lower grain units, while higher sediment loads corresponds to higher grain units in general. This means that if the measures are not implemented, there will be higher crop yield (grain units) but there will also be higher sediment load. There is also a trade off between grain units and maximum flow. As the measures are implemented, lower maximum flow corresponds to lower grain units, while higher maximum flow corresponds to higher grain units. The same is observed between frequency of low flow days and grain units.

Cluster Results

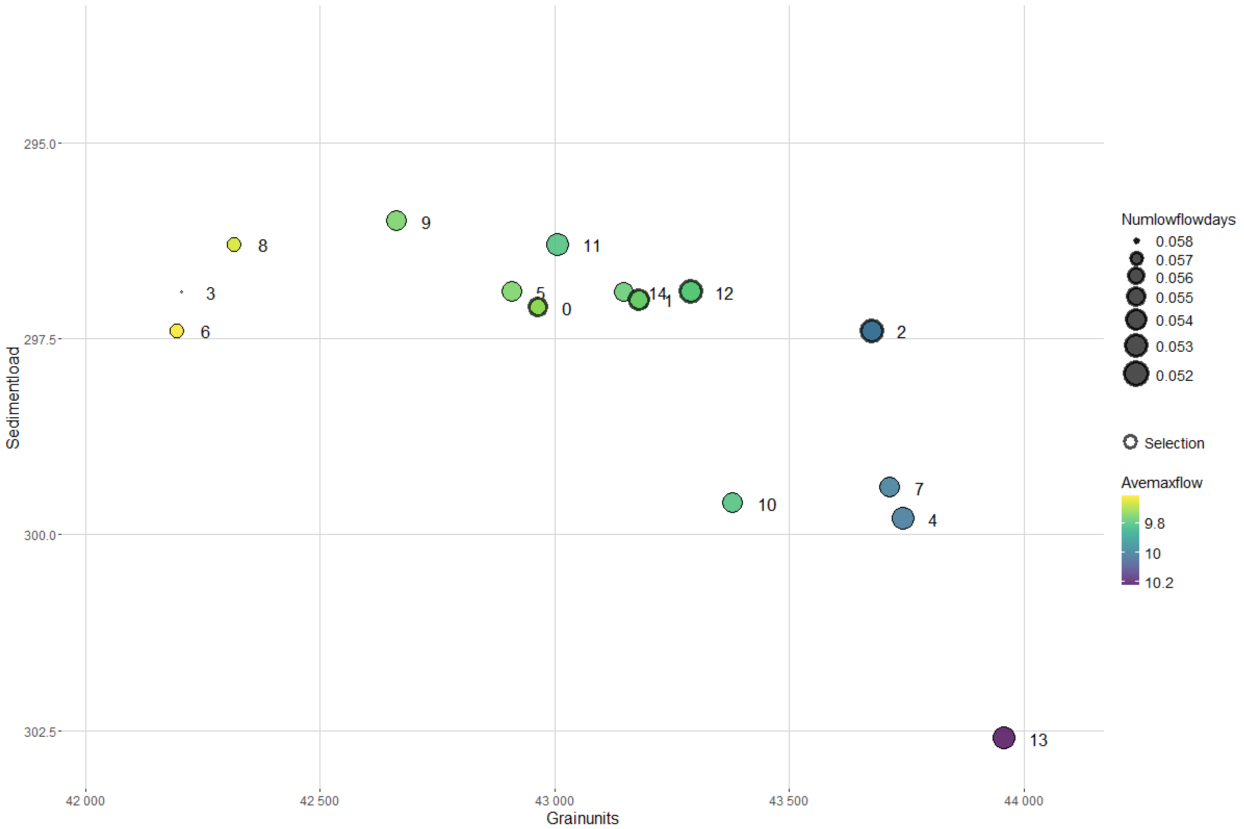

Figure A15: Scatter plot of the Pareto optima remaining after clustering. The three largest clusters (0, 1, 2, 12) are marked in bold (from left to right). Example from case study 7 Wimbe.

The clustering, performed with default settings, produced 15 well-separated clusters (Figure A15). The original Pareto front shape remains visible and well represented by the remaining optima. The three largest clusters are nr. 0, 1 and 2, with cluster 2 having a cluster size of 113. Cluster 13 achieves highest grain unit performance, and it also resulted in the highest sediment load and maximum average flow.

Cluster Analysis

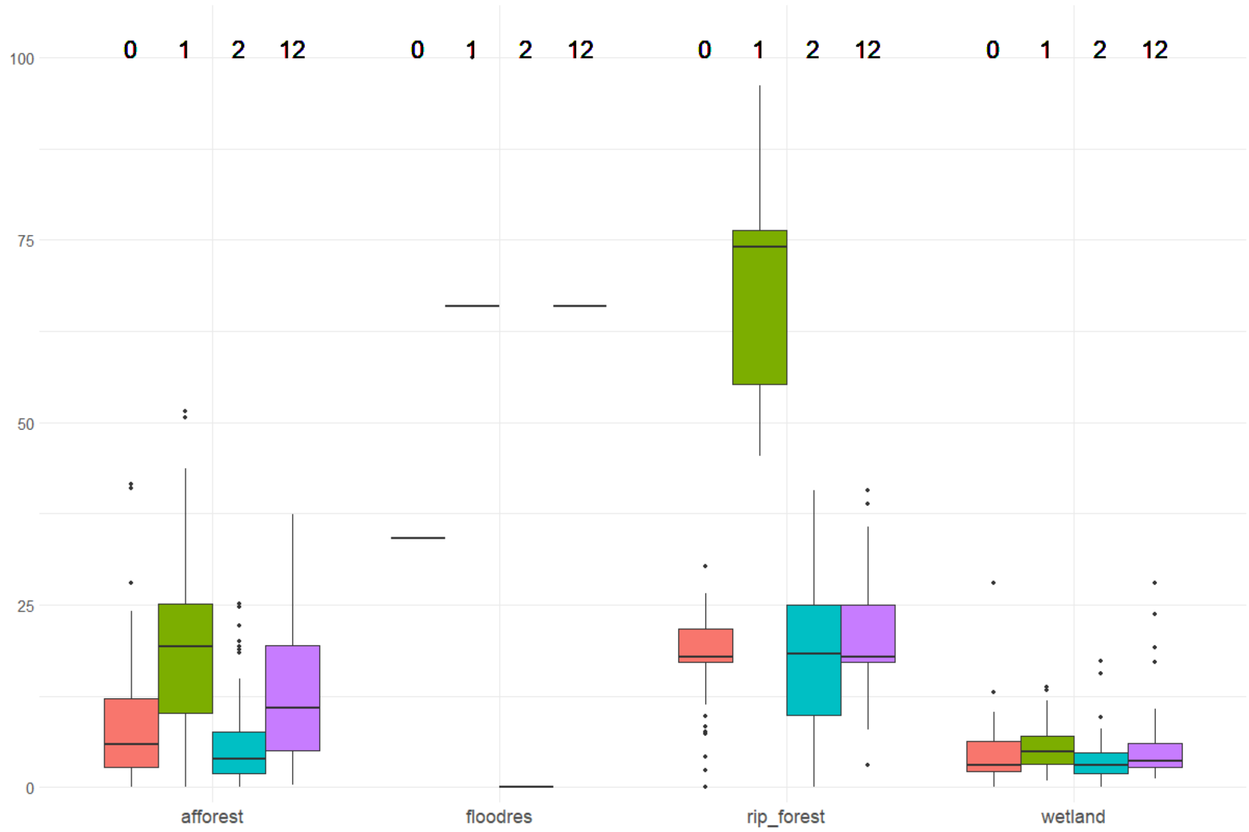

Figure A 16: The individual measures' share in total considered area. Example clusters 0, 1, 2 and 12) from case study 7 Wimbe.

As shown in Figure A 16, Cluster 1 has the highest proportion of riparian forest while all 4 clusters (0, 1, 2 and 12) have almost the same proportions of wetlands. The proportion of flood restoration measures are the same in clusters 1 and 12. The different clusters suggest more afforestation efforts either in the riparian zone or within the field areas. The other measures, riparian buffers and hedges, are not yet shown as technical fixes are still to be implemented.

A8 - Case Study 8 - Dotnuvele, Lithuania

The SWAT+ model setup for CS 8 was finished on 14 December 2024 (project partner - KU). The team also analysed the effectiveness of NSWRMs using the calibrated SWAT+ model and quantified them for four scenarios: riparian buffer strips, cover crops, reduce tillage, and small wetlands.

The tasks required for the MOO (running the scenarios using the SWATmeasR package, optimisation using CoMOLA, then applying the post-processing using ParetoPick-R) are currently being performed. As the lead modeller is in demand for the main modelling tool (SWAT+) development and shares the responsibility for the CS, there have been some delays in the completion of modelling tasks for the Lithuanian CS. The work is scheduled to be completed by the end of May 2025.

A9 - Case Study 9 - Cherio, Italy

<u>Full Pareto front</u>

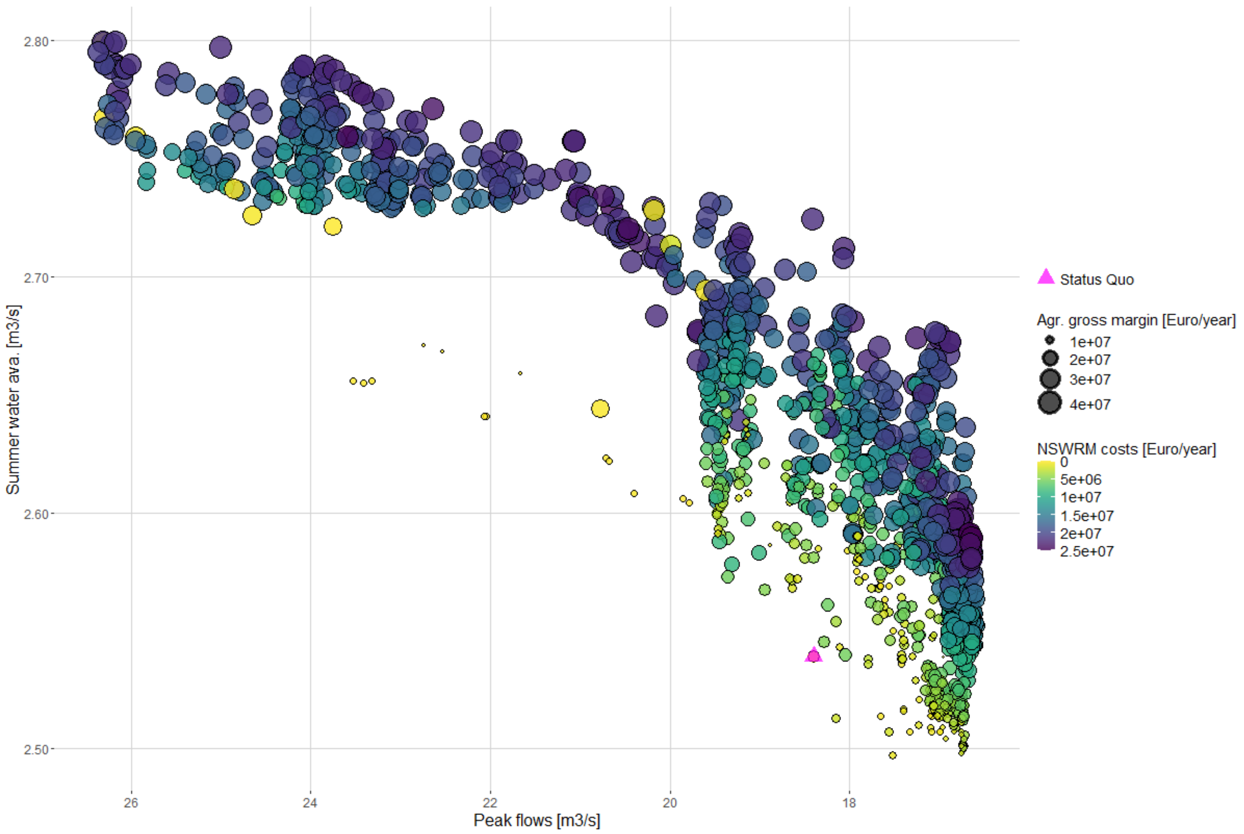

Figure A17: Scatter plot of the full Pareto front. Example from case study 9 Cherio River basin.

The x and y axes represent hydrological objectives. The Pareto front demonstrates how certain scenarios that enhance water availability to levels significant for the agricultural sector during the irrigation season are simultaneously effective in mitigating peak flow intensity. Generally, the most expensive scenarios tend to maximize benefits for hydrological objectives, but significantly less costly alternatives appear to perform nearly as well. The relatively large size and darker colour of the dots on the outer front suggest that a high level of measure implementation does not necessarily hinder agricultural production.

Frequency Analysis

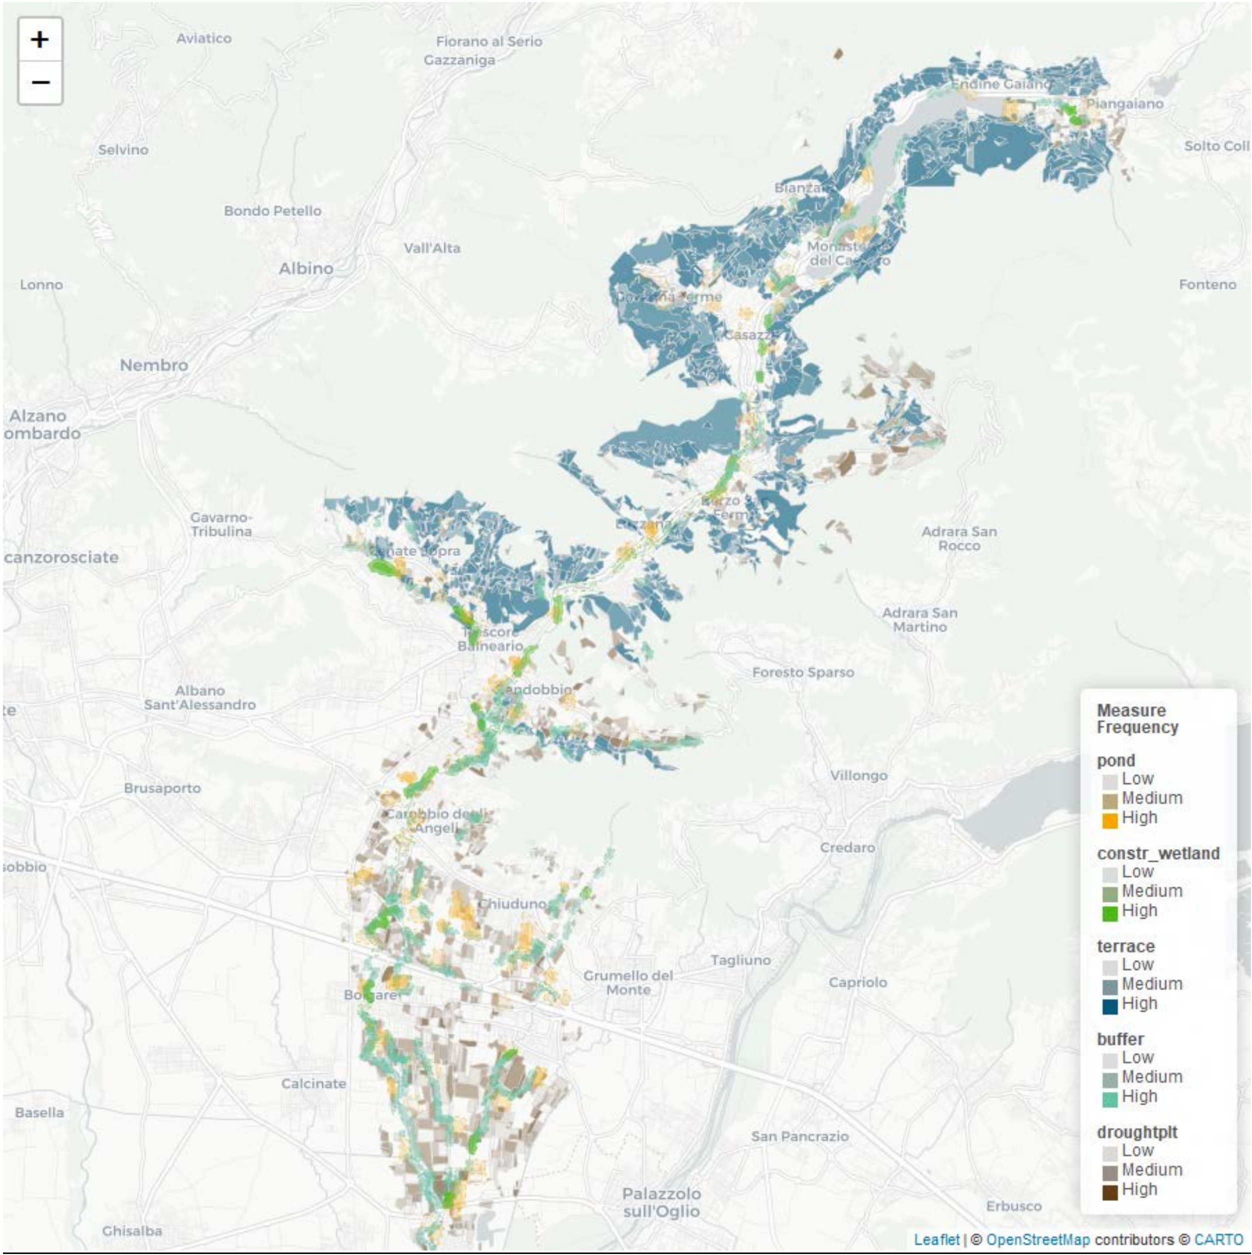

Figure A18: Summary of all measure implementation plans. Frequency with which individual HRUs are activated in those optima performing beyond 0.5 across all objectives. Example from CS 9 - Cherio River basin.

Constructed wetlands are the most frequently implemented measures in the Italian CS, likely due to their strong positive impact on environmental objectives and relatively low implementation costs. Ponds are more commonly implemented in the lower portion of the basin, where the need for stream overflow prevention is greater. The implementation frequency of terraces is medium-high; their single drawback, high construction costs, is well counter-balanced by their beneficial effects on hydrological resilience and the ability to enhance agricultural productivity in the areas where they are applied.

Cluster Results

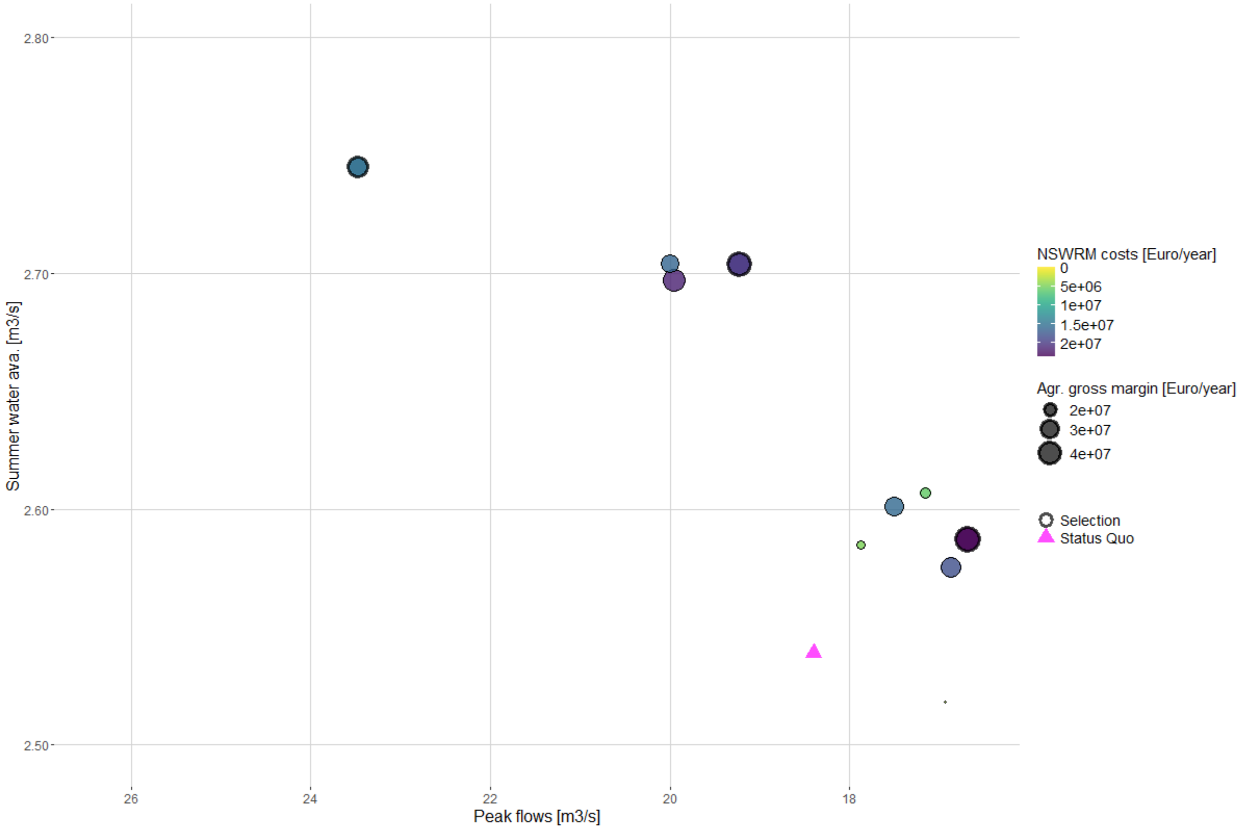

Figure A19: Scatter plot of the Pareto optima remaining after clustering. Three clusters have been marked in bold (from left to right: 9, 0, 6). Example from case study 9 Cherio River basin.

Using the k-means method, 10 clusters were generated for the Italian CS. Cluster 9 represents scenarios aimed at maximizing water availability during the summer, characterised by medium implementation costs. Cluster 0 represents the optimal compromise between the two hydrological objectives, although it is associated with high costs. Cluster 6, along with adjacent clusters, illustrates how peak flows can be mitigated at varying implementation costs, each with differing consequences for agricultural production.

Cluster Analysis

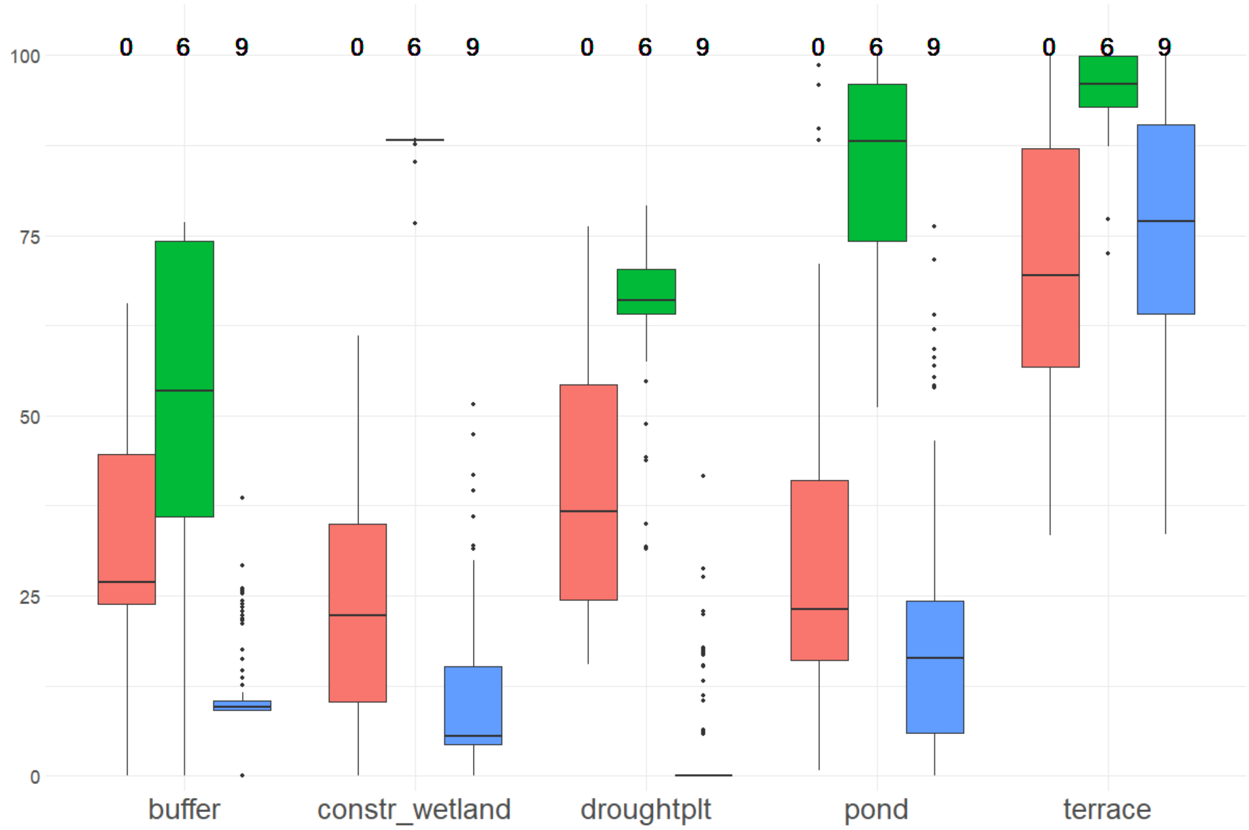

Figure A20: The individual measures' share in total considered area. Example clusters 0, 9 and 6) from case study 9 Cherio river basin.

The most consistently implemented measure across the three selected clusters is terraces. In contrast, the implementation of other measures, such as ponds or drought-resistant plants, exhibits greater variability. This may be due to the fact that these measures are highly effective for benefitting specific hydrological processes but could be counterproductive for others.

A10 - Case Study 10 - Kråkstad, Norway

Full Pareto front

The following four indicators were used for preliminary MOO of the Kråkstadelva catchment in Norway: 1) Q_minmax: ratio between maximum and minimum flow (Q_max/Q_min) at the catchment outlet; 2) Tot N_load: total N load; 3) N_conc days: frequency of days when the total N concentration is below threshold, corresponding to good water quality according to the Water Framework Directive and 4) crop yield, expressed in average annual grain units. For the final optimisation, implementation costs will replace one of the N objectives.

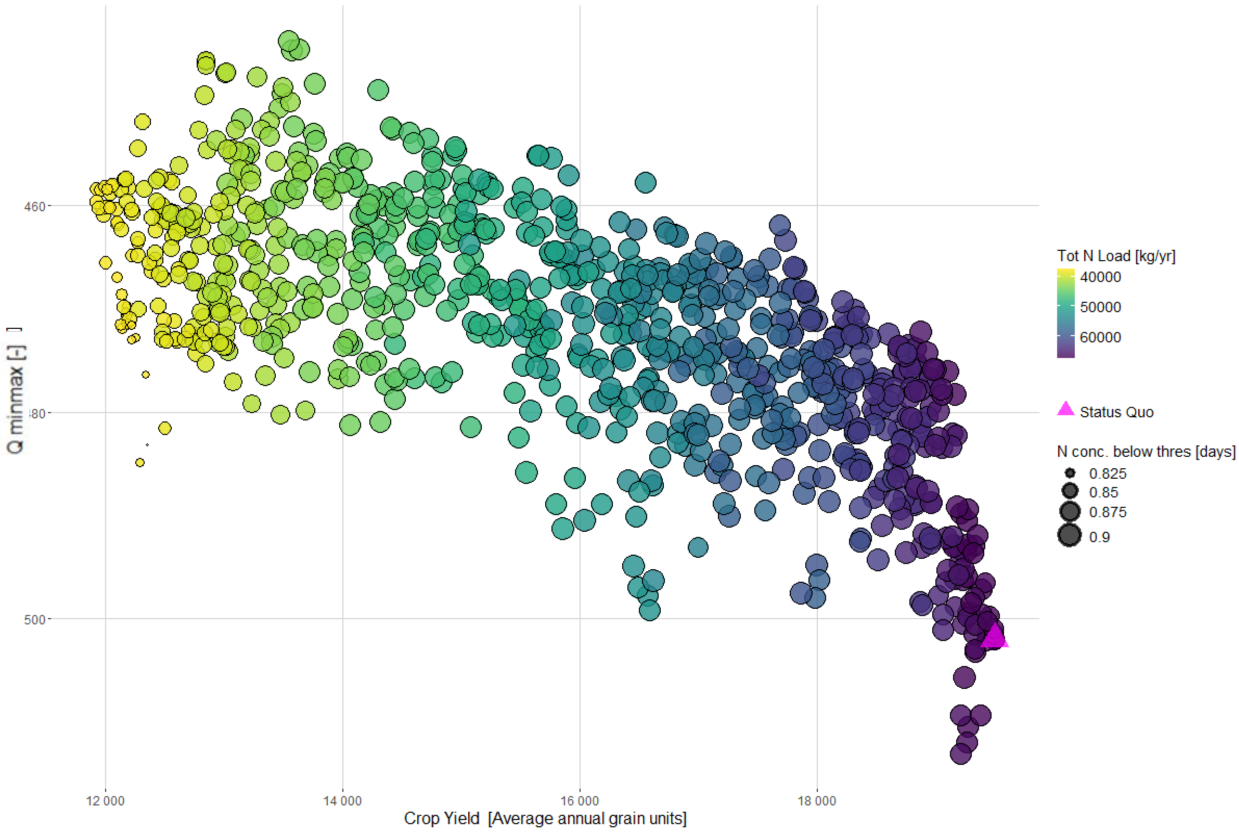

Figure A21: Scatter plot of the full Pareto front. Example from case study 10 Kråkstad.

The catchment representation during the reference period (status quo, 2015-2018) is characterised by relatively low levels of measure implementation, resulting in high N loads, high yields, low water retention and low N concentrations. The Pareto front (Figure A21) indicates that some improvement to water retention could be made without sacrificing yields, but that strong reductions in N loads must be met with yield losses. Reductions in Q_minmax combined with reductions in N loads cannot be achieved without increases in N concentrations, which is likely due to the dilution effect, as higher water retention in the landscape results in reduced flow, resulting in higher total N concentrations in the stream even with reduced N loads. This effect, however, needs further investigation. A high implementation level of both, management and structural measures would have a positive impact on water retention and flow regulation (expressed in reduction of Q_max/Q_min ratio by approx. 45), and would lead to significant reduction in N loads (from more than 60000 kg/year to less than 40000 kg/year).

Frequency Analysis

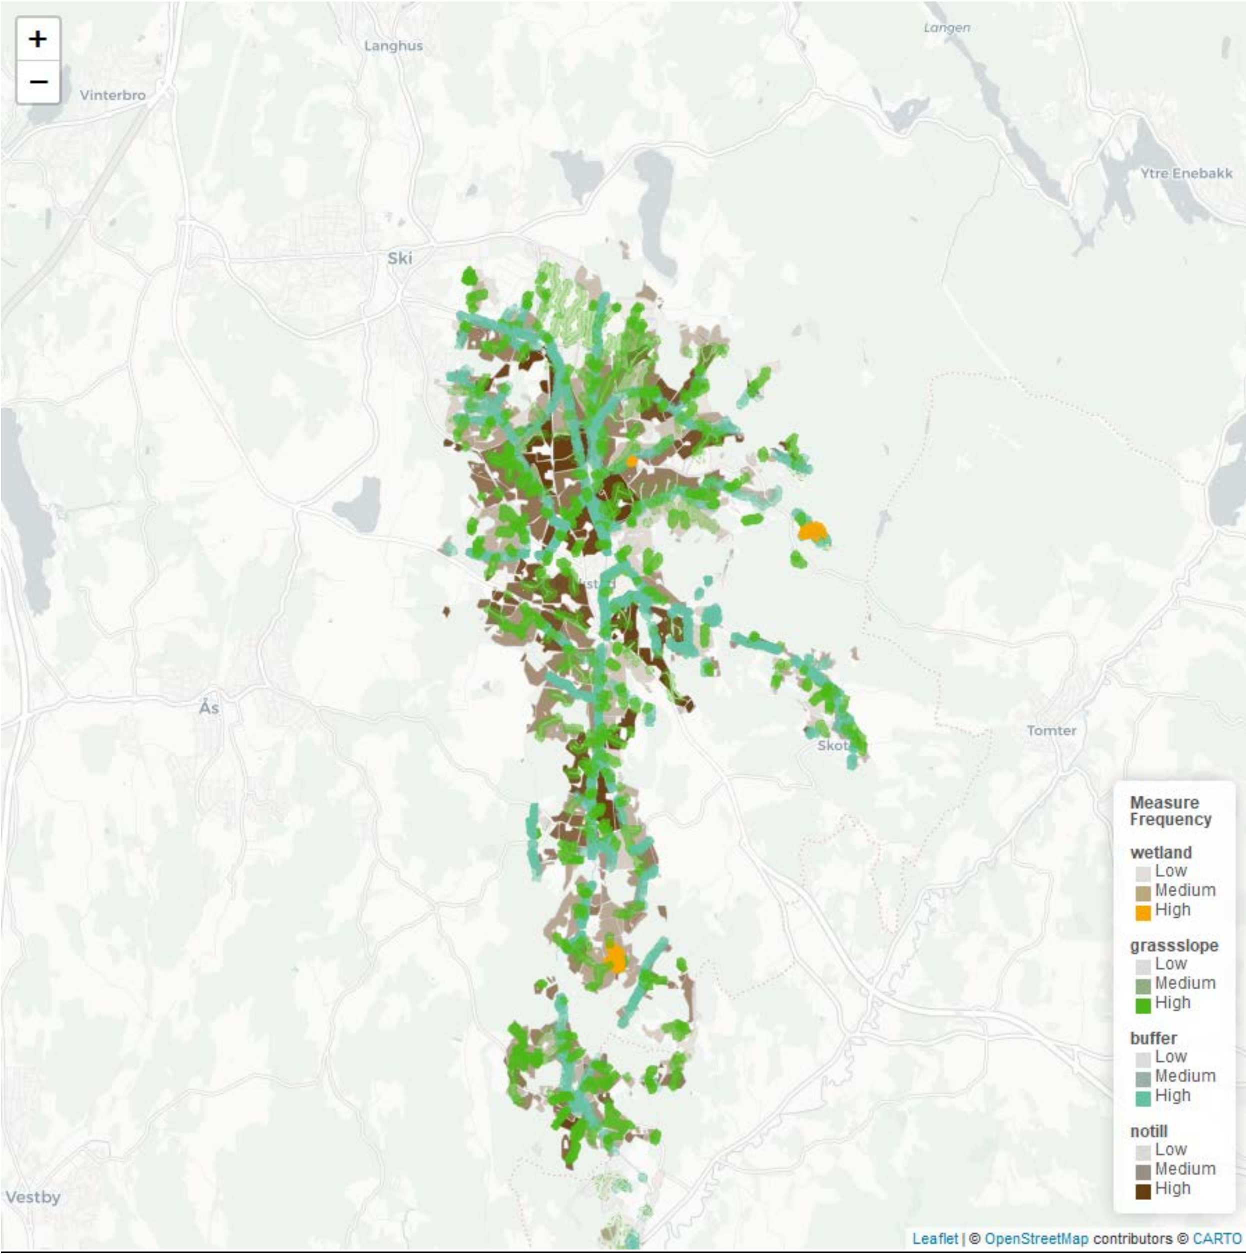

Figure A22: Summary of all measure implementation plans. Frequency with which individual HRUs are activated in those optima performing beyond 0.5 across all objectives.

Figure A22 shows an example frequency plot for all optima with above-average (>0.5 scaled) performance across all objectives. The orange patterns indicate three constructed wetlands that are implemented with high frequency. Conservation tillage (consisting of no tillage in the autumn and stubble during the winter period and in early spring) is implemented with high or moderate frequency for most of the agricultural fields. The distance between the fields and "*the closest stream*" seems to be more important than the erosion risk class of the individual fields for frequency of implementation. Grassed waterways appeared to be the second most important measure for flow and nutrient loads regulation after lowtill.

For high water retention and low N loads, notill is frequently implemented. Solutions with higher yields avoid notill and rely strongly on buffer strips, and to a lesser extent grassed waterways. Low N concentration solutions employ a medium level notill, complemented by frequent buffer strips and grassed waterways. Constructed wetlands are commonly implemented across all solutions.

Cluster Results

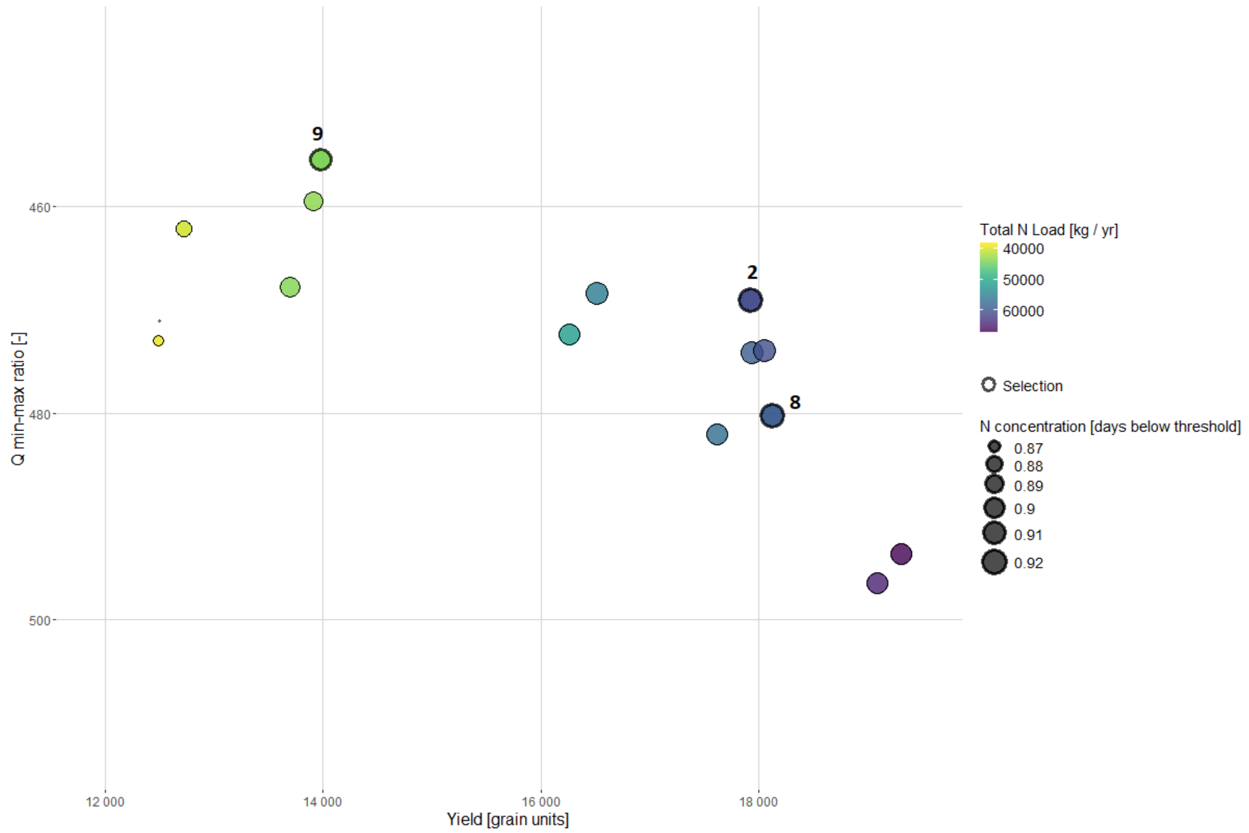

Figure A23: Scatter plot of the Pareto optima remaining after clustering. The three largest clusters (9, 2, 8) are marked in bold (from left to right). Example from case study 10 Kråkstad

The clustering has been performed with default settings. A more in-depth tuning will be performed on the final results. A clustering of 15 is shown in Figure A23, fairly well spaced out and still representative of the general shape of the pareto front.

Cluster Analysis

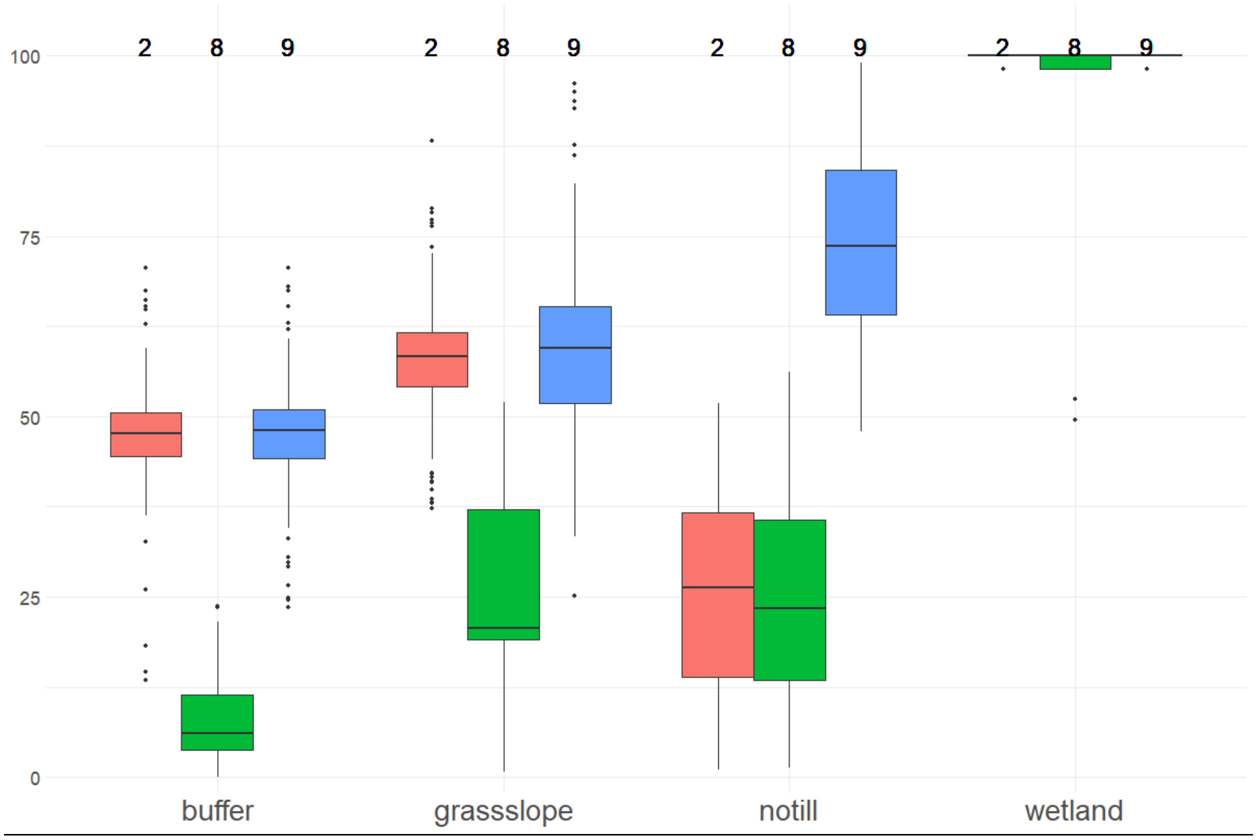

Figure A24: The individual measures' share in total considered area. Example clusters 9, 2, and 8 from case study 10 - Kråkstad.

Across the 3 largest clusters, wetlands find high degrees of implementation (Figure A24). Cluster 9, with the highest water retention, has the highest implementation of measures in general, but specifically notill. Cluster 2, finding a good balance between yield and water retention, has lower implementation of notill, but retains the high implementation rates of grass slopes and buffers of cluster 9. Cluster 8 matches cluster 2 in notill, but has a lower implementation rate for grass slopes and especially buffer strips. This leads to a worse water retention with little improvement to yield, N loads or N concentrations. Cluster 8 highlights the need for a more fine-tuned clustering approach.

Outlook

In the future, we plan to replace grain units with gross margin; however, the gross margin calculation has not been finalised. Moreover, we might consider using other indicators, including sediment yields and/or frequency of days when the total P concentration is below the Water Framework Directive threshold. The results, included in this report, together with those from the new calculations will be presented to the stakeholders during the 3rd MARG meeting.

A11 - Case Study 11 - Tetves, Hungary

For the analyses of the optimal spatial allocation of the NSWRMs in CS 11, the following environmental and economic indicators were applied for the optimisation computations: N load in the channel (kg/a), average soil moisture content of the top 30 cm of the soil at the agricultural parcels of the catchment (mm), grain unit of the agricultural parcels of the catchment (unit/a), implementation and maintenance costs of the NSWRMs (Euro/a).

<u>Full Pareto front</u>

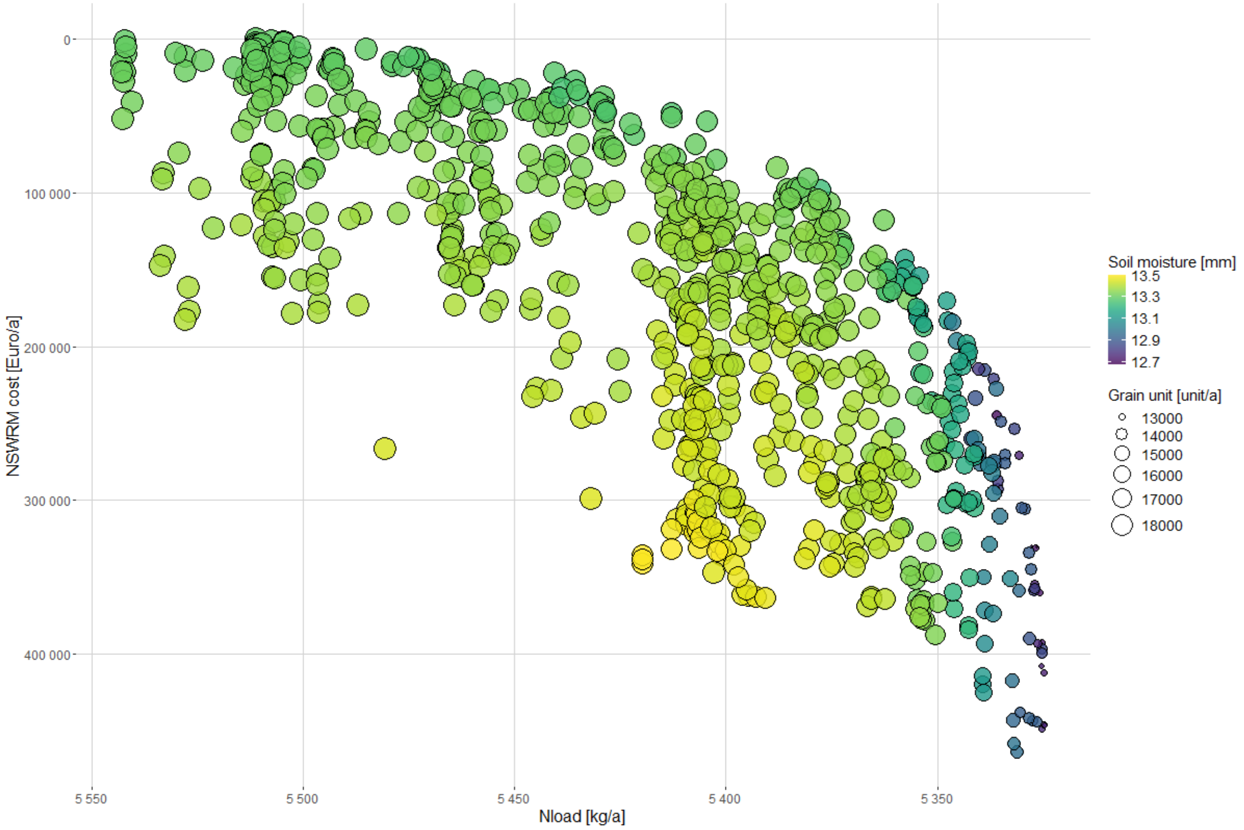

Figure A25: Scatter plot of the full Pareto front. Example from case study 11 Tetves.

The Pareto front of the Tetves study area (Figure A25) illustrates the relationship between the N load in the stream, total grain yields of croplands, soil moisture content in the top 30 cm, and the implementation cost of NSWRMs. The figure highlights that, in general, yields decline when measures to reduce N loss and/or soil moisture loss are implemented. However, it is also evident that good crop yields can still be achieved while maintaining soil moisture and N loss at acceptable levels, albeit with increased expenditure.

Frequency Analysis

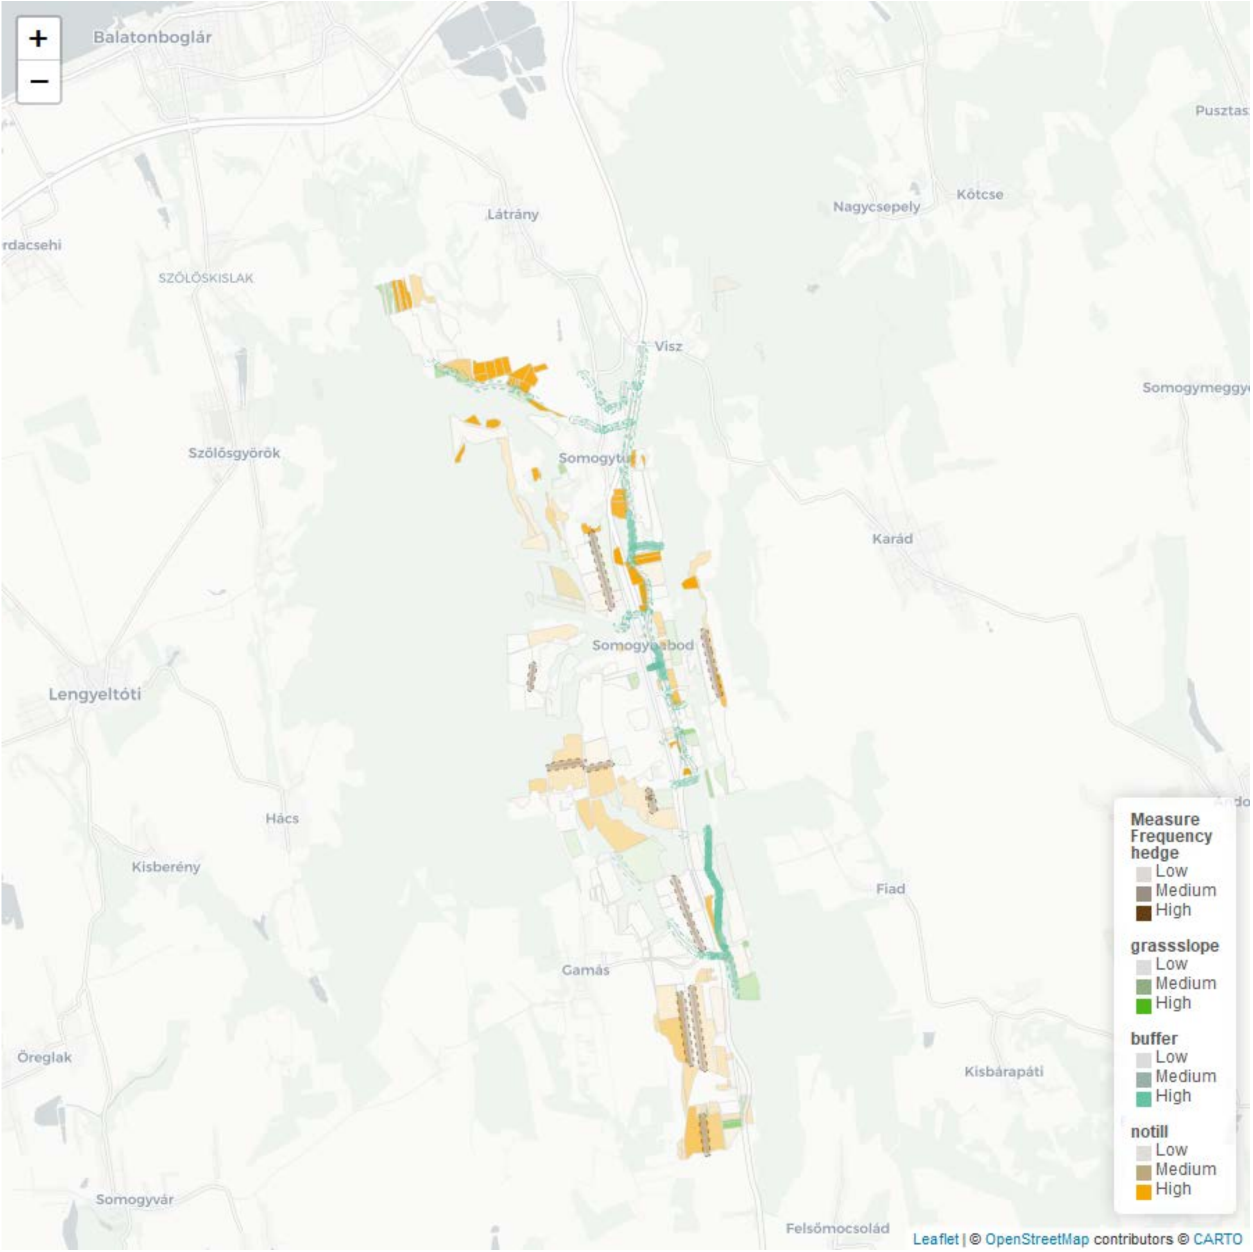

Figure A26: Summary of all measure implementation plans. Frequency with which individual HRUs are activated in those optima performing beyond 0.5 across all objectives. Example from case study 11 Tetves.

The overall measure frequencies range from medium to low across all measures, with some concentrated high-frequency hotspots for riparian buffer strips, no-till management with cover crops, and field-dividing hedges in areas where erosive processes are severe (Figure A26).

Cluster Results

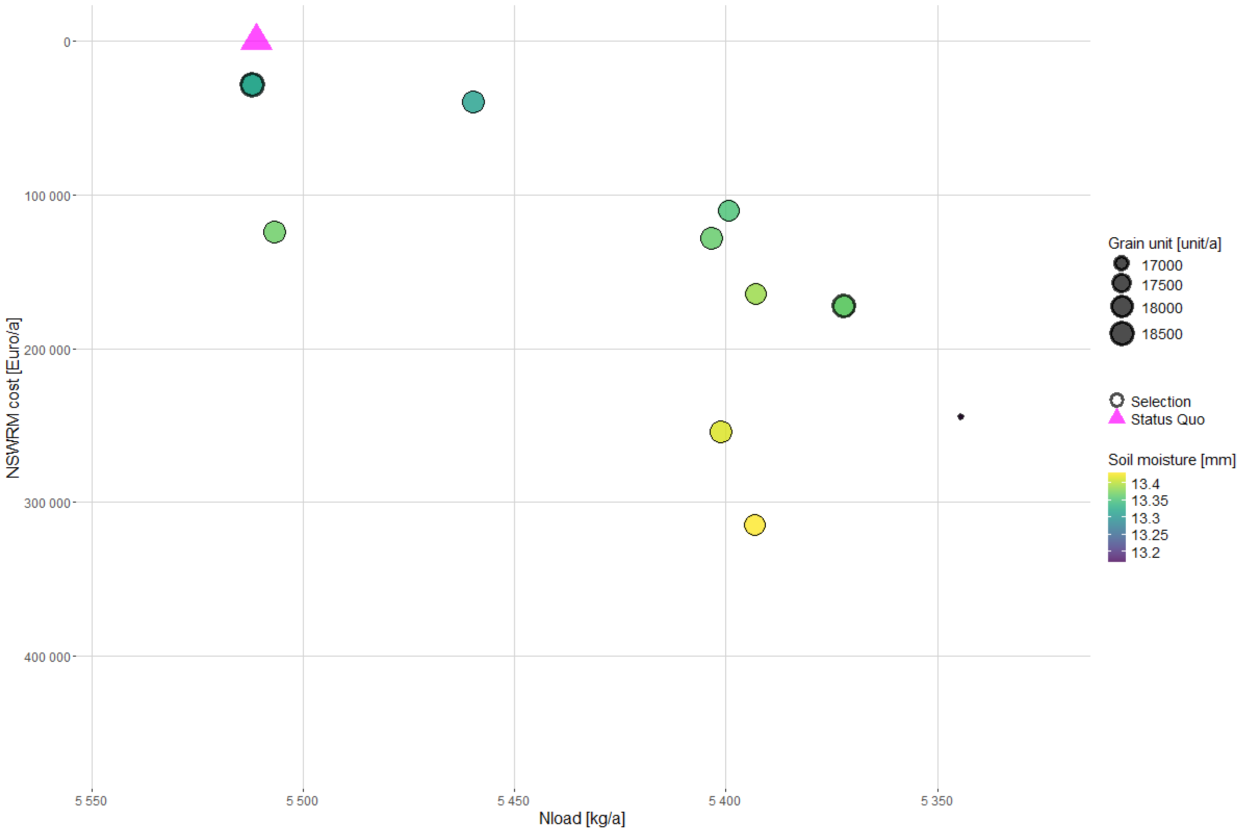

Figure A27: Scatter plot of the Pareto optima remaining after clustering. The three largest clusters (9, 3, 1) are marked in bold (from left to right). Example from case study 11 Tetves.

The k-medoid clustering (with a preliminarily set minimum 2 and maximum 20 cluster number) delineated ten separate clusters (Figure A27). The three largest clusters are highly differentiated, positioned at both ends and the center of the Pareto front. Cluster 1 is characterised by low grain yield, a significant reduction in N load (mean: 5350 kg/a), and high implementation costs. Cluster 3 exhibits moderate values across all four analyzed indicators, while Cluster 9 closely resembles the status quo, with unchanged N load, medium soil moisture, high grain yield, and low NSWRM costs.

Cluster Analysis

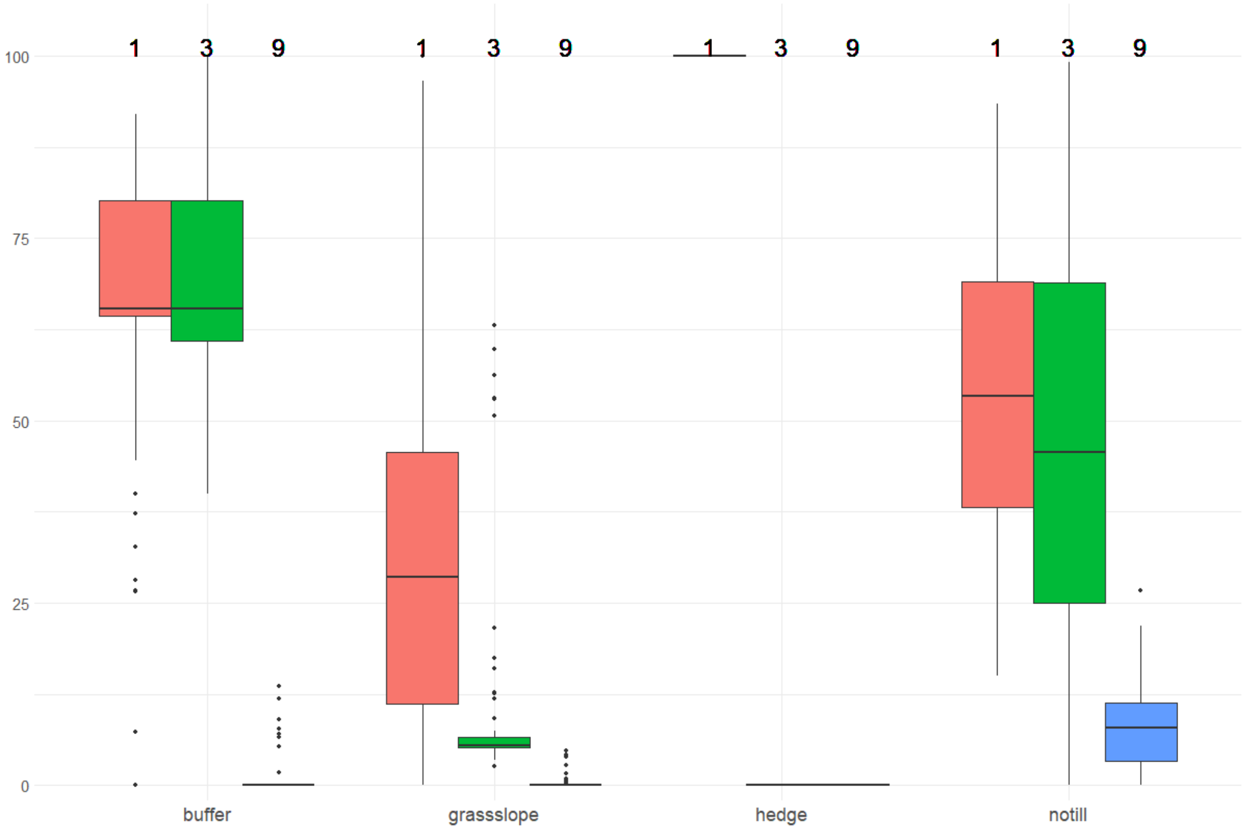

Figure A28: The individual measures' share in total considered area. Example clusters 1, 3 and 9 from case study 11 - Tetves.

The boxplot diagram of the three largest clusters (1, 3 and 9) highlights the key differences between these three scenarios (Figure A28). Cluster 1 features extensive implementation of riparian buffers, no-till management with cover crops, and arable land conversion to grassland (grass slopes). Additionally, this cluster has all the hedges activated (all hedges were switched on or off together during the optimisation phase). The proportion of converting arable land to grassland causes a strong decrease in grain unit of the main crops of the area and lower soil moisture content in the top 30 cm. The number of activated different types of measures resulted in a scenario with high NSWRM implementation costs. Cluster 3 primarily includes riparian buffers and no-till management with cover crops, while Custer 9 almost avoids all of the measures except no-till management on a few agricultural fields.

Outlook

For the presentation of the optimisation results at the MARG interviews, NSWRM costs shall be replaced by gross margin, as this economic indicator might be more informative for farmers and farm advisors.

A12 - Case Study 12 - Cechticky, Czechia

CS 12 - Čechtický catchment is not yet at a stage where it is possible to present results that would correspond to the expected content of this annex. Work is currently underway on the Common Optimisation Protocol (D5.1). The setup of the SWATmeasR project has been completed. For this purpose, the following three steps were performed: (1) initializing a new SWATmeasR project, (2) retrieving the definitions (parameterisations) of all NSWRMs to be implemented in the model setup, and (3) defining the location of all NSWRMs to be implemented.

Also, the selection of indicators relevant for the construction of the scatter plot of the full Pareto front was performed:

- average annual sum of grain units in whole basin
- frequency daily discharge is below high flow threshold
- average annual N loss from land objects
- NSWRM implementation costs

Attention will now focus on the implementation of measures in the CoMOLA workflow and on the MOO run. Once these have been completed, they can be evaluated in post processing.

A13 - Case Study 13 - Dviete, Latvia

The team responsible for this CS did not provide updates on their progress in time to be included in this deliverable. At the time of the most recent update from the CS, the modelling had not yet progressed to the optimisation stage.

A14 - Case Study 14 - Sävjaån, Sweden

The team responsible for this CS did not provide updates on their progress in time to be included in this deliverable. At the time of the most recent update from the CS, the modelling had not yet progressed to the optimisation stage.