




Fuzzy MCDM

User Guide

How to install Fuzzy MCDM and carry out a criteria-weighting study

Version 1.0.0

2026-10-15

Ing. Akram Elsaied Elomiya

ORCID <https://orcid.org/0000-0002-7166-0239> University of Pardubice, Faculty of Transport Engineering, Department of Transport Management, Marketing, and Logistics, Studentská 95, 532 10 Pardubice, Czech Republic

Tanta University, Faculty of Engineering, Department of Public Works Engineering, 31511, Tanta, Egypt

Supervisor: doc. Ing. Jiří Křupka, Ph.D.

This software is supported by the project SGS_2026_008 'Modelling of Transport Technology and Management Selected Aspects VI' at the University of Pardubice.

Fuzzy MCDM 1.0.0

© 2026 University of Pardubice. Distributed under the Fuzzy MCDM Research Licence 1.0.

Website: <https://fuzzy-mcdm.pages.dev>

DOI: [10.5281/zenodo.23380643](https://doi.org/10.5281/zenodo.23380643)

This software is supported by the project SGS_2026_008 'Modelling of Transport Technology and Management Selected Aspects VI' at the University of Pardubice.

Microsoft and Windows are trademarks of the Microsoft group of companies. Other names may be trademarks of their respective owners.

Contents

1	Introduction	1
1.1	What the application does	1
1.2	How a study proceeds	1
1.3	Conventions	2
1.4	Citing the software	2
2	Installation	3
2.1	System requirements	3
2.2	Download and verify	3
2.3	Install	3
2.4	First start	3
2.5	Update and uninstall	4
2.6	Where your files are	4
3	Your first study	5
3.1	The Welcome screen	5
3.2	Creating a study with the setup assistant	5
3.3	The Overview screen	6
3.4	Refining criteria and experts afterwards	6
3.5	Saving your study	7
4	Entering judgments	9
4.1	Entering a pair on the selector	9
4.2	Saving, skipping and drafts	9
4.3	The comparison matrix	10
4.4	Scale & mapping	10
4.5	Importing judgments	10
4.5.1	Input formats	10
4.5.2	Signed levels	10
4.5.3	Problems reported	10
5	Consistency, methods and the run	13
5.1	Quality review and the 0.10 rule	13
5.2	Choosing methods and comparators	13
5.3	Pre-run review and starting the run	14
6	Reading the results	17
6.1	The Results pages and the run bar	17
6.2	Summary	17
6.3	Weights and rankings	17
6.4	Method comparison	18
6.5	Consistency and expert disagreement	19
6.6	Leave-one-expert-out	19
6.7	Resampling stability	20
6.8	The figure pages	20
6.9	Next step	20

7	Evidence and export	23
7.1	Inspect calculations	23
7.2	Calculation trace	23
7.3	References	23
7.4	Provenance	23
7.5	The Export page	24
7.6	Exporting a single table	25
8	Study files and settings	27
8.1	Study files	27
8.2	File compatibility	27
8.3	Recovery copies	27
8.4	Changes made by other programs	27
8.5	Settings	27
8.6	Native menu	28
8.7	Where your files are	28
9	Troubleshooting	29
A	Example study	31
A.1	Study	31
A.2	Judgments	31
A.3	Run	32
	References	33

Chapter 1

Introduction

This guide explains how to install Fuzzy MCDM and how to carry out a criteria-weighting study with it, from the first criterion to the exported report. It describes version 1.0.0.

1.1 What the application does

Fuzzy MCDM determines the weights of decision criteria from pairwise judgments of several experts. You describe the decision as a goal with main criteria and sub-criteria, each expert compares the criteria in pairs on Saaty's 1–9 scale, and the application computes the weights with five weighting methods from the same judgments:

- AHP with the eigenvector method (Saaty, 1980);
- fuzzy AHP with the geometric mean method (Buckley, 1985);
- fuzzy AHP with the fuzzy inverse of column sum method, FICSM (Ahmed & Kilic, 2019);
- interval type-2 fuzzy AHP (Kahraman et al., 2014) with the full 1–9 scale of Beresneva and Gordenko (2018);
- interval type-2 fuzzy AHP with the scale of Milošević et al. (2019), kept as a historical reference.

Two sets of equal weights, one over the hierarchy and one flat, are shown beside the methods as comparators. They are not weighting methods.

One consistency rule applies to every method: a judgment matrix of an expert whose consistency ratio is above 0.10 is excluded from the primary result. Because all methods use the same judgments and the same rule, the differences between their weights come from the methods alone. The application shows those differences, how much the experts disagree, how the weights change when one expert is left out, and how stable they are under resampling of the experts.

Every number can be traced back. The calculation trace shows each step from the judgments to the final weights, with the equation, the values substituted into it and the publication that defines it. Each run carries fingerprints of its inputs and outputs, so a result in a report can be matched to the run that produced it.

Fuzzy MCDM runs fully offline. It has no network functions, sends no data and has no automatic updates. Your studies stay in the files you save.

1.2 How a study proceeds

A study follows the order of the sidebar:

1. create the study: name, goal, criteria and experts (Chapter 3);
2. enter or import the experts' pairwise judgments (Chapter 4);
3. review consistency, choose the methods and analyses, and run (Chapter 5);

4. read the results and figures (Chapter 6);
5. inspect the calculations and export tables, figures, a report or a reproducibility bundle (Chapter 7).

Chapter 8 describes study files, recovery and settings, Chapter 9 lists common problems, and Appendix A gives a complete synthetic example study that you can open and repeat.

1.3 Conventions

Names of screens, buttons and menu items appear as **Pairwise Comparisons**. File names appear as `study.mcdmlab`. Menu paths are written `File > Save`. The screenshots were taken with the synthetic example study of Appendix A; the data in them are not real expert judgments.

1.4 Citing the software

If you publish results obtained with Fuzzy MCDM, cite the software and the version you used:

Elomiya, A. E. (2026). *Fuzzy MCDM* (Version 1.0.0) [Computer software]. University of Pardubice. <https://fuzzy-mcdm.pages.dev>

The website also provides the citation in BibTeX and CFF formats.

Chapter 2

Installation

This chapter explains how to install and set up Fuzzy MCDM on a Windows computer.

2.1 System requirements

Fuzzy MCDM runs on Windows 10 or Windows 11, 64-bit (x64). The installer works for the current user only; administrator rights are not required.

2.2 Download and verify

Download the installer `Fuzzy MCDM-1.0.0-win-x64.exe` from the <https://fuzzy-mcdm.pages.dev> website (Download section). The same page shows its SHA-256 checksum.

Verify the checksum in PowerShell:

```
Get-FileHash -Algorithm SHA256 '.\Fuzzy MCDM-<version>-win-x64.exe'
```

Replace `<version>` with the version you downloaded. The printed hash must match the published one. If it differs, delete the file and download it again.

The installer is not code-signed. Windows SmartScreen may show **Windows protected your PC**. Choose **More info**, check that the file name is `Fuzzy MCDM-1.0.0-win-x64.exe`, then select **Run anyway**.

2.3 Install

Run the installer. Follow these steps:

1. Read and accept the **Fuzzy MCDM Research Licence** on the licence agreement page.
2. Choose to install for the current user.
3. Select the installation folder (the default is `%LOCALAPPDATA%\Programs\Fuzzy MCDM`).
4. Click **Install**.
5. Click **Finish** when the process completes.

A **Fuzzy MCDM** entry appears in the Start menu and a desktop shortcut is created.

2.4 First start

Open Fuzzy MCDM from the Start menu or the desktop shortcut. The application opens on the Welcome screen. It works offline and does not require any network access.

2.5 Update and uninstall

Fuzzy MCDM has no automatic updates. To update, download and run the newer installer; existing study files are kept.

To uninstall: go to **Windows Settings > Apps > Installed apps**, find Fuzzy MCDM, and select **Uninstall**. Study files are not removed.

2.6 Where your files are

- Study files (.mcdmlab) are saved wherever you choose when you create or export them.
- Settings and automatic recovery copies are stored in %APPDATA%\@mcdm\desktop.

For more details, see Chapter 8.

Chapter 3

Your first study

This chapter shows how to create your first study: from the Welcome screen, through the four-step setup assistant, to the Overview screen, where you can refine the criteria and experts afterwards, and how saving works.

3.1 The Welcome screen

After installation (Chapter 2), the application opens on the Welcome screen. It shows the **Fuzzy MCDM** wordmark and the thesis attribution, the tiles **New Study** and **Open Study**, and a **Recent studies** table with the columns Name, Location and Modified. When autosaved work exists, a banner such as **2 studies have unsaved work — Review** appears; it opens the recovery copies (Chapter 8). A theme selector and an **About** link are also available.




Figure 3.1. The Welcome screen.

3.2 Creating a study with the setup assistant

Start a new study from the tile or with File > New Study (Ctrl+N). The setup assistant opens as a full window with four steps:

1. **Study:** enter the **Study name** and the **Decision goal**.
2. **Criteria:** add the main criteria; they receive the codes C1, C2, and so on. Use **Add subcriteria to <name>** to add sub-criteria with the codes C1.1, C1.2, and so on, and **Paste a list** to add many names at once. Duplicate names are refused.

3. **Experts:** add the experts; each row receives a code such as E1, E2, and so on. **More experts** adds a row.
4. **Methods & review:** tick the methods for the study. The comparators are listed for orientation. Reasons that block the creation are listed on the screen. Confirm with **Create study**.

New study Cancel

Criteria
Main criteria are compared for the decision goal. Subcriteria are optional and are compared only with their siblings under the same parent. A study with only main criteria is complete.

Main criteria - 3 + > Paste a list

C1 Cost Remove

Subcriteria of Cost
- 2 +

C1.1 Purchase cost Remove

C1.2 Operating cost Remove

[Remove subcriteria](#)

C2 Service quality Remove

Subcriteria of Service quality
- 2 +

C2.1 Reliability Remove

C2.2 Passenger comfort Remove

[Remove subcriteria](#)

C3 Environment Remove

Subcriteria of Environment
- 2 +

Back Next

Figure 3.2. The Criteria step of the setup assistant.

Only verified methods can be ticked. You can change the selection later on the **Methods** page.

If you leave the assistant with **Cancel**, the application asks **Discard this new study?**.

3.3 The Overview screen

After **Create study**, the application opens the Overview screen. It shows the study title, the decision goal and the project details (the file path, or **Not saved to a file yet**). A **Next step** card carries a button that opens the next screen, and a **Progress** panel shows how far the study has advanced.

The sidebar is the quickest way to move through a study. Its permanent groups are **STUDY** (**Overview**, **Criteria**, **Experts**, **Methods**), **JUDGMENTS** (**Pairwise Comparisons**, **Quality Review**), **ANALYSIS** (**Pre-run Review**) and **RESULTS** (**Results**), plus **Settings**.

3.4 Refining criteria and experts afterwards

You can change the structure on the Criteria screen. Rename the decision goal, use **Add main criterion** to add a main criterion, and open the **More** menu on a row for **Rename**, **Add subcri-**




Figure 3.3. The Overview screen.

terion, Delete and Move (Up, Down, Indent, Outdent). If a change would invalidate entered judgments or stored results, a dialog explains the impact before anything is changed.

On the Experts screen, each row holds one expert with a display label and an influence weight: any finite number greater than or equal to zero. The normalised share β_e is shown next to it. You can add an expert, remove one (the application asks for confirmation) and reset all experts to equal influence.

3.5 Saving your study

Save with File > Save (Ctrl+S) or File > Save As (Ctrl+Shift+S). Studies are stored as .mcdmlab files, wherever you choose. The header shows the save status: **Saved, Unsaved changes • recovery copy kept** or **Not saved to a file • recovery copy kept**. Autosave keeps recovery copies in the background; the interval is set in **Settings**.

Judgments are entered on Pairwise Comparisons, covered in Chapter 4.

Chapter 4

Entering judgments

This chapter explains how to enter pairwise judgments on Pairwise Comparisons: how to read the 17-position selector, how saving, skipping and drafts work, how the comparison matrix tracks progress, what the Scale & mapping screen shows, and how to import judgments from a file or from text.

4.1 Entering a pair on the selector

On Pairwise Comparisons, choose the **Comparison set** and the **Expert**. A comparison set is the set of main criteria, or the sub-criteria of one main criterion. Each pair is entered on a 17-position selector. Positions 9 to 2 on the left give the row criterion, the centre position 1 means **Equal**, and positions 2 to 9 on the right give the column criterion. For example, “row 4” means the row criterion is 4 times as important.

Urban bus fleet (synthetic example) Not saved to a file · recovery copy kept Save Save As... Next: Methods

Pairwise Comparisons Import judgments Scale & mapping Continue to quality review

Comparison set Main criteria — 3/3 · complete Expert Expert B — 6/6

Comparing main criteria for: Weight the criteria for choosing an urban bus fleet technology

This set: 3 of 3 comparisons Expert: 6 of 6 comparisons Study: 18 of 18 comparisons

	C1	C2	C3
C1	1	1/2 ✓	2 ✓
C2	2	1	3 ✓
C3	1/2	1/3	1

C1 Cost - C2 Service quality - C3 Environment

C1 Cost vs C3 Environment Comparison 2 of 3 in this set Saved

C1 Cost C3 Environment

9	8	7	6	5	4	3	2	1	2	3	4	5	6	7	8	9
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

Cost is between equally and moderately more important than Environment (2)

> Scale help Saaty 1-9 levels

Skip for now No changes to save Save Save — all comparisons entered

Figure 4.1. Entering a judgment on the 17-position selector.

4.2 Saving, skipping and drafts

Confirm a pair with **Save & Next**; on the last pair the button reads **Save — all comparisons entered**. **Skip** leaves the pair without saving. If you open a pair that already holds a saved judgment and your draft differs from it, the application asks you to choose **Use saved judgment** or **Keep my draft**. If you leave the screen with an unsaved draft, you are asked to **Stay**, **Discard** or **Save**.

4.3 The comparison matrix

The list **Comparisons in this set** and a comparison matrix show how much of the set is already entered. Clicking a matrix cell opens that pair, so you can return to an earlier pair at any time.

4.4 Scale & mapping

The header button **Scale & mapping** opens a read-only view of the Saaty 1-9 scale and its verbal labels (Saaty, 1980). The screen also has a **Technical details** toggle. Level 1 means equal.

4.5 Importing judgments

Instead of entering every pair by hand, you can import judgments from a file or from text. The header button **Import judgments** opens a panel in Pairwise Comparisons with six stages: **Select**, **Detect**, **Map**, **Preview**, **Validate** and **Commit**.

On **Select** you choose the source: **From file** (CSV, TSV, TXT or XLSX), or **Paste text** with **Paste from clipboard** and **Import this text**.

On **Map** you configure how the source becomes judgments. Choose the **Mapping mode** (**Long format** or **Matrix format**) and the columns for **Expert**, **Parent**, **Row criterion**, **Column criterion** and **Value**. You map unknown verbal terms, resolve duplicates with **Keep existing** or **Use imported**, and accept the exclusions. An option keeps the raw file in the study.

Nothing is repaired silently. **Commit** stays disabled while conflicts are unresolved, exclusions are not accepted, or rows are ambiguous.

4.5.1 Input formats

The long table has one row per comparison, as CSV or TSV. Its columns are, for example:

```
expert,parent,row_criterion,column_criterion,value
```

A square matrix CSV carries a header row such as , c1, c2, c3 and the row labels in the first column; reciprocal conflicts are flagged. An XLSX file holds one long-table sheet, or one square matrix per sheet; cells with formulas or dates are refused. When you paste, use tab- or comma-separated text, as a matrix or as a long table.

Every cell holds one value: a level from 1 to 9, a fraction such as 1/3, a decimal equal to a scale value, a scale label such as **Moderate importance**, or a verbal term that you map during the import.

4.5.2 Signed levels

Tick **Values are signed levels (3 or +3 means 3, -3 means 1/3)** when the source stores intensities as signed levels: one optional sign and one digit from 1 to 9 per cell, so -3 means 1/3. The values 0 and -1 are refused.

Signed levels are not allowed in XLSX. Save the workbook as CSV first.

4.5.3 Problems reported

The import reports problems instead of repairing anything silently. Typical problems are:

- a missing judgment,

- a duplicate (identical or contradictory),
- an invalid intensity,
- an invalid reciprocal,
- a reciprocity conflict,
- an unknown criterion,
- an unknown expert,
- an unmapped verbal term,
- a pair across different parents,
- an incomplete matrix,
- malformed input.

The header button **Continue to quality review** takes you on to the consistency check described in Chapter 5.

Chapter 5

Consistency, methods and the run

This chapter explains how consistency is checked in Quality Review, how to choose the weighting methods and comparators for a run, and how to configure the options and start the run from Pre-run Review.

5.1 Quality review and the 0.10 rule

Because each expert enters their own judgments (see Chapter 4), consistency is checked separately for every expert and comparison set. Quality Review shows one card per expert and comparison set, each with Saaty's consistency ratio (CR). The threshold is 0.10: an expert's matrix with CR above 0.10 is excluded from the primary result for that comparison set, while the expert's other matrices stay in. Sets of two criteria are always consistent (CR = 0).




Figure 5.1. The Quality Review screen.

A matrix with a consistency ratio above 0.10 is excluded from the primary result. The same rule applies to every method, so the methods are compared on the same judgments.

5.2 Choosing methods and comparators

On the Methods screen you tick the weighting methods and comparators for the run. The five weighting methods are:

- AHP with the eigenvector method (Saaty, 1980);

- fuzzy AHP with the geometric mean method (Buckley, 1985);
- fuzzy AHP with the fuzzy inverse of column sum method (FICSM) (Ahmed & Kilic, 2019);
- interval type-2 fuzzy AHP with the full 1-9 scale (Beresneva & Gordenko, 2018; Kahraman et al., 2014);
- interval type-2 fuzzy AHP with the scale of Milošević et al. (2019), kept as a historical reference.

Beside them, the comparators, equal weights over the hierarchy and flat equal weights, are listed for orientation; they are not weighting methods. You can select all verified methods at once. A checkbox **Leave-one-expert-out** adds a run over the experts; it needs at least two experts.




Figure 5.2. The Methods screen.

View formulas and sources opens a side panel per method with the equations behind the method and their sources. Check it to see exactly what a method computes before you run it.

5.3 Pre-run review and starting the run

The run starts from **Pre-run Review**. The screen shows a **Study summary**, an **Analysis options** block, the list of expert influences and the **Analysis availability**.

In **Analysis options** you confirm the comparators and **Leave-one-out**, and you can tick **Bootstrap (resampling stability)** with **Replicates (B)**, a whole number from 10 to 10 000, and **Seed**, a whole number from 0 to $2^{64} - 1$; **New seed** draws another seed. The same seed and number of replicates give the same results again.

Analysis availability tells you which analyses can run. Run blockers are listed below it, each with a button that opens the screen that fixes it. If the study changed since the last visit, **Pre-run review is out of date** appears, with a reload. Resolve the blockers, then start the run from this screen.




Figure 5.3. The Pre-run Review screen.

Leave-one-out needs at least two experts.

Reading the results

This chapter explains how to read a run: which pages the results are split into, what the weights and rankings mean, how to check consistency and stability, and how to set the figures up for printing.

6.1 The Results pages and the run bar

Open **Results** in the sidebar to reach the output of a run. The left rail becomes a tree with one entry per page:

- **Summary**
- **Weights & rankings**, holding **Global weights**, **Local weights** and **Rankings**
- **Method comparison**
- **Quality**, holding **Consistency** and **Expert disagreement**
- **Sensitivity**, holding **Leave-one-expert-out** and **Resampling stability**
- **Figures**, holding five figure pages
- **Evidence**, holding **Inspect calculations**, **Calculation trace**, **References** and **Provenance**
- **Export**

Above the page, the run bar tells you which run you are looking at. It carries the run selector, a run status, the run ID, the input revision and the output hash. A run marked **Current** was produced from the judgments the study holds now. Runs marked **Historical** are earlier runs kept with the study, and runs marked **Legacy** come from an older version of the application. The run ID and the output hash are the identifiers you quote in correspondence about a result.

To read an earlier run, switch to it:

1. Select **Change run** in the run bar.
2. Pick the run you want from the list.

Every page below follows the run you have selected.

6.2 Summary

The **Summary** page opens first. It gives you the final global weight of every leaf criterion, the criterion that ranks first, and how far the methods agree with one another. Read it as a check on the run before you look at the detail.

6.3 Weights and rankings

The **Weights & rankings** group holds three tables.

- **Global weights** are the final weights of the leaf criteria. They are the numbers that answer the decision goal.




Figure 6.1. The Summary page of a run, with the final global weights, the top-ranked criterion and the method agreement.

- **Local weights** are the weights inside one parent, before they are combined with the weight of that parent. Use them to see how a main criterion is split among its sub-criteria.
- **Rankings** list the criteria in order, so that you can compare methods by their ordering rather than by their numbers.

All the table pages share one toolbar:

- **Methods** shows how many of the selected methods are on screen, as N of M. Reduce it to compare two methods instead of five.
- The criteria scope and the comparison set narrow the table to part of the hierarchy.
- The search box keeps only the criteria whose name contains what you type.
- **Inspect** opens the calculation behind a value.

6.4 Method comparison

The **Method comparison** page places the selected methods side by side, so that you can see where they agree and where they differ. The study can compute criteria weights with five methods:

- AHP by the eigenvector method (Saaty, 1980);
- fuzzy AHP with the geometric mean (Buckley, 1985);
- fuzzy AHP by the fuzzy inverse of column sum method, FICSM (Ahmed & Kilic, 2019);
- interval type-2 fuzzy AHP on the full 1–9 scale (Beresneva & Gordenko, 2018; Kahraman et al., 2014);




Figure 6.2. The Global weights page, with the shared table toolbar above the weights of every leaf criterion.

- interval type-2 fuzzy AHP with the scale of Milošević et al. (2019), kept as a historical reference.

Two comparators can be shown beside them: **equal weights (hierarchical)** and **equal weights (flat)**. They are not weighting methods. They show how far the judgments move the weights away from equal weights.

6.5 Consistency and expert disagreement

The **Consistency** page repeats the consistency ratio (CR) of each expert matrix, and the **Expert disagreement** page shows how far the experts differ from one another. Read both before you quote a ranking.

The rules are the same for every method:

- The threshold is a CR of 0.10.
- A matrix with a CR above 0.10 is excluded from the primary result.
- A comparison set of two criteria is always consistent, so its CR is 0.

A run that excludes a matrix is still complete. Check which matrices were kept before you use the weights, because an excluded matrix removes one expert from one comparison set and the remaining experts carry the set alone.

6.6 Leave-one-expert-out

The **Leave-one-expert-out** page shows how much one expert decides. It repeats the run once for every expert, each time without that expert, and shows how the weights move. An expert whose removal changes the weights strongly has a decisive influence on the result.

The analysis needs at least two experts; with one expert the option blocks the run on the **Pre-run Review** page.

6.7 Resampling stability

The **Resampling stability** page shows how much the result depends on the particular judgments. It resamples the experts with replacement many times, recomputes the weights, and reports a range for each weight between the 5th and the 95th percentile. A narrow range means the weights hold whichever experts are drawn; a wide range means they depend on individual experts. The same number of replicates and the same seed give the same ranges.

Set it up on the **Pre-run Review** screen, not on the results page:

1. Tick **Bootstrap (resampling stability)**.
2. Give a whole number from 10 to 10000 in **Replicates (B)**.
3. Give a whole number from 0 to $2^{64} - 1$ in **Seed**, or select **New seed**.
4. Run the study again.

6.8 The figure pages

The **Figures** group holds five pages: **Weights by method**, **Ranks by method**, **Expert judgments**, **Leave-one-out influence** and **Resampling ranges**. Each page shows one chart and an inspector beside it.

Above the chart, the toolbar offers:

- **Fit** and **Actual size** to change the zoom;
- **Undo** and **Redo** to step through your changes;
- **Reset to preset** to go back to the settings of the current preset;
- **Show values as table** to read the numbers instead of the chart;
- **Export...** to save the chart as a file.

The inspector beside the chart has four tabs, **Data**, **Appearance**, **Labels** and **Output**. Together they set:

- the chart type: vertical bars, horizontal bars, dot or slope;
- a preset: **Screen**, **Report**, **Journal** or **Presentation**;
- the palette: accessible, greyscale, or mono with hatching;
- the background, font sizes, titles, axes, legend, data labels and the number of decimals.

Figure settings are saved with the study, so a colleague who opens the file sees the same charts.

6.9 Next step

Save the study when the run is finished, so that its inputs, results, trace and provenance travel with the file. Then select **Export** in the results tree to write the evidence out. Chapter 7 describes what each export contains.




Figure 6.3. The Weights by method figure page, with the chart toolbar above it and the inspector on the right.

Evidence and export

This chapter explains how you check a single weight against the underlying data, how the application records where every number came from, and what each export writes to disk.

7.1 Inspect calculations

A weight in a results table is the end of a chain of steps. When you want to see the part of the chain that produced one value, select **Inspect**. The application opens the calculation for that value instead of the whole table, so you read one chain and not a page of numbers. The full chain, for every value in the run, is on the **Calculation trace** page.

7.2 Calculation trace

The **Calculation trace** page lists every step from the judgments to the weights. Each step gives you the equation that was used, the values that were substituted into it, and the source the equation comes from. Nothing is summarised away, so you can follow one weight back to the judgments an expert entered.

The **Final result** selector at the top chooses which result you are tracing, for example the local weights of the main criteria under one method. The steps then run from that result down to the inputs, in the order the computation uses them: the matrix built from the judgments, the consistency ratio of that matrix, whether the matrix passed, how the expert matrices were aggregated into one, and the priority vector that is the result itself.

To follow a weight:

1. Select **Calculation trace** under **Evidence** in the results tree.
2. Choose the result in **Final result**.
3. Read the steps from the top. Each one opens with **Technical details** if you want the equation and the substituted values.

The matrices in the trace are shown in signed form, the form in which the judgments are stored: a value s above 1 means that the row criterion is s times as important, and a negative value $-s$ means that the column criterion is s times as important, that is the ratio $1/s$. The value 1 means equal importance.

7.3 References

The **References** page lists the publications behind the methods the run used. It is the same list that the **Bibliography** export writes, so a reader of your report can go to the source of every method without leaving the document.

7.4 Provenance

The **Provenance** page records what the run was made with and what it produced:




Figure 7.1. The Calculation trace page, showing the steps from the priority vector down to the matrix an expert entered.

- the application version;
- the method profiles that were used;
- fingerprints of the inputs and of the outputs.

A fingerprint is a checksum of one piece of data. The input fingerprints tell you that the judgments are the ones you intended; the output fingerprint tells you that the weights in a document are the ones the run produced. Quote them when you report a result, because they change as soon as the run changes.

7.5 The Export page

Select **Export** at the end of the results tree. The page offers five exports. They write separate files, so you can make as many as you need.

1. **Tables & weights (calculation book)** writes one Excel workbook with every table at full stored precision. Use it when the numbers will be reused or checked in a spreadsheet.
2. **Figures** writes each figure as SVG or PNG. The resolution is 300 dpi by default, and the background is white or transparent as you choose.
3. **Report** writes a printable A4 PDF, the **Calculation report**. It contains the study and the run, the inputs and which responses were included, the final weights, the method comparison, consistency, sensitivity, the sources, and the provenance identifiers.
4. **Reproducibility (run bundle)** writes one zip archive. It holds the project, the judgments, the validation results, the method profiles, the weights, the comparisons, the calculation book, the figures, the tables as CSV, Markdown and LaTeX, the provenance, and a MANIFEST.sha256 list of checksums.




Figure 7.2. The Provenance page, with the application version, the method profiles and the input and output fingerprints.

5. **Bibliography** writes the references of the methods used, as BibTeX, TXT, Markdown or PDF.

The **Report** export is not available for a legacy run that has no input snapshot, because the report names the inputs and there are none to name.

When a matrix was excluded because its CR was above the threshold, a **Diagnostic run (includes CR-failing matrices)** JSON is offered as well. It is a second set of weights that keeps the failing matrices. Use it to see what the excluded judgment would have changed, not as the result of the study.

7.6 Exporting a single table

Every table has its own export menu, so you do not have to take the whole calculation book to obtain one table. The menu offers CSV, XLSX, Markdown and LaTeX.

1. Open the table and find its export menu.
2. Choose the format.
3. Save the file where you want it.

The format decides how much evidence travels with the table:

- A CSV is accompanied by a `<name>_provenance.txt` file written beside it.
- An XLSX carries a **Provenance** sheet.
- Markdown and LaTeX carry no extra file, so cite the run ID and the output hash from the run bar if you use them.




Figure 7.3. The Export page, with the five exports offered for the current run.

A single table exported as CSV or XLSX keeps its provenance, but it does not keep the judgments. Take the run bundle as well if the reader has to follow the numbers back to the input.

When the study is finished, select File > Save to keep the run, its trace and its provenance in the .mcdmlab file. Chapter 6 describes what the saved pages show.

Chapter 8

Study files and settings

This chapter explains how studies are saved, how recovery works, and how to configure the application.

8.1 Study files

A study is saved as a `.mcdmlab` file. This is a ZIP archive that contains the project, all judgments, every run with its inputs and results, and a `MANIFEST.sha256` checksum list. On open, the checksums are verified. If any entry does not match, the study opens read-only and the mismatched entries are listed. `File > Save` is then disabled.

Study files can be saved anywhere on your computer. The file location appears in the Overview page.

8.2 File compatibility

Files saved with versions 0.1.0 and 0.2.0 open. Criteria, experts, and judgments are preserved. Stored runs made with methods that were later removed from the application are dropped on open, and a notice names them. A converted study shows **Converted from an older format — save to update** in the header. When you save, a `.bak` copy of the old file is kept.

Files from newer versions are refused.

8.3 Recovery copies

Autosave keeps recovery copies in the background. The interval is configurable in Settings (30 to 600 seconds, default 120). After an unexpected close, the Welcome screen offers to review unsaved work.

8.4 Changes made by other programs

If the study file was modified by another program since it was opened, saving asks first, so that the other changes are not overwritten without warning.

Only one copy of the application runs at a time. Starting it again brings the open window to the front.

8.5 Settings

The Settings screen allows you to configure:

- **Autosave interval:** Seconds between automatic recovery saves (30 to 600).
- **Retain raw files by default:** Whether imported source files are kept in the study.
- **Appearance:** System, Light, or Dark theme.

To open Settings, select **Settings** in the left sidebar. After making changes, click **Save**.

8.6 Native menu

The native menu provides file and application controls:

- **File:** **New Study** (Ctrl+N), **Open Study** (Ctrl+O), **Open Recent**, **Save** (Ctrl+S), **Save As** (Ctrl+Shift+S), **Close Study** (Ctrl+W), **Exit**.
- **Edit:** Standard editing commands.
- **View:** Zoom controls, full screen, **Appearance** (System, Light, Dark).
- **Help:** **About Fuzzy MCDM**.

8.7 Where your files are

- Study files (.mcdmlab) are saved wherever you choose when you create or export them.
- Settings, recent files, and recovery copies are stored in %APPDATA%\@mcdm\desktop.

Exports are described in Chapter 7.

Troubleshooting

This chapter lists common problems and how to resolve them.

Checksum mismatch on open If a `.mcdmLab` file fails verification, the study opens read-only and lists the mismatched entries. You cannot save it. Open an earlier copy of the file if you have one.

Dropped runs notice When opening a study from an older version, runs that use methods removed from the application are dropped. A notice lists the removed methods. The study itself opens normally with all criteria, experts, and judgments preserved.

Newer-format file refused Files from future application versions cannot be opened. Ensure you are running the latest version from <https://fuzzy-mcdm.pages.dev>.

Import problems When importing judgments, the following problems are reported and must be resolved before committing:

- Missing judgment
- Duplicate (identical or contradictory)
- Invalid intensity or invalid reciprocal
- Reciprocity conflict
- Unknown criterion or unknown expert
- Unmapped verbal term
- Pair across different parents
- Incomplete matrix
- Malformed input

Nothing is repaired silently. Resolve conflicts, accept exclusions, and map unknown terms before the import can be committed.

Run blockers The **Pre-run Review** page lists any conditions that prevent a run. Each blocker has a button that opens the relevant screen to fix it.

Leave-one-out needs two experts The leave-one-expert-out analysis requires at least two experts. Add another expert on the **Experts** page, or clear the **Leave-one-expert-out** option.

CR above 0.10 excluded Judgments with a consistency ratio (CR) above 0.10 are excluded from the primary result. This applies to all weighting methods. The **Quality Review** page shows which matrices are excluded. Ask the expert to revise the judgments, or accept the exclusion.

Changed-on-disk prompt If the study file was modified by another program since it was opened, saving asks first to avoid overwriting external changes.

Recovery copies After an unexpected close, the Welcome screen offers to review recovery copies. These are autosaved states from unsaved work. To adjust the autosave interval, open Settings and change the interval (30 to 600 seconds).

Study files and recovery are described in Chapter 8; installer warnings in Chapter 2.

Appendix A

Example study

This appendix gives the complete synthetic study used for the screenshots in this guide. You can download it from the website as `urban-bus-fleet-example.mcdmlab` and open it with File > Open Study, or enter it yourself to practise every step. The judgments are invented for illustration; they are not the opinions of real experts.

A.1 Study

Study name Urban bus fleet (synthetic example)

Decision goal Weight the criteria for choosing an urban bus fleet technology

Experts Expert A, Expert B and Expert C, with equal influence

Table A.1. Criteria of the example study.

Main criterion	Sub-criteria
C1 Cost	C1.1 Purchase cost, C1.2 Operating cost
C2 Service quality	C2.1 Reliability, C2.2 Passenger comfort
C3 Environment	C3.1 Emissions, C3.2 Noise

A.2 Judgments

Table A.2 lists every judgment. Each pair is written with its row criterion first. “Row 4” means that the row criterion is four times as important as the column criterion: on the selector, choose 4 on the left side. “Column 2” means that the column criterion is twice as important: choose 2 on the right side. “Equal” is the centre position.

Table A.2. Pairwise judgments of the example study.

Pair	Expert A	Expert B	Expert C
Cost vs Service quality	row 2	column 2	row 3
Cost vs Environment	row 4	row 2	row 3
Service quality vs Environment	row 2	row 3	equal
Purchase cost vs Operating cost	row 2	column 2	column 3
Reliability vs Passenger comfort	row 3	row 2	row 5
Emissions vs Noise	row 4	row 2	row 3

All matrices pass the consistency check. The consistency ratios of the main-criteria matrices are 0.0000 (Expert A), 0.0079 (Expert B) and 0.0000 (Expert C). A set of two criteria is always consistent, so the sub-criteria matrices have a ratio of zero.

A.3 Run

The screenshots use all five weighting methods, leave-one-expert-out and the bootstrap with 500 replicates and seed 20261013. The bootstrap ranges therefore repeat exactly when you run the study with the same seed.

Table A.3 gives the resulting global weights of the sub-criteria, rounded to four decimals. The application shows and exports them at full precision.

Table A.3. Global weights of the sub-criteria in the example study.

Sub-criterion	AHP	Fuzzy AHP (Buckley)	Fuzzy AHP (FICSM)	IT2 fuzzy AHP (full 1–9)	IT2 fuzzy AHP (historical)
Purchase cost	0.2018	0.2019	0.1997	0.2045	0.2019
Operating cost	0.2910	0.2873	0.2806	0.2790	0.2882
Reliability	0.2503	0.2459	0.2487	0.2453	0.2470
Passenger comfort	0.0806	0.0833	0.0860	0.0859	0.0826
Emissions	0.1310	0.1330	0.1347	0.1341	0.1324
Noise	0.0454	0.0486	0.0503	0.0511	0.0478

All five methods rank Operating cost first and Noise last. The methods differ most for Operating cost, where the weights range from 0.2790 to 0.2910. The **Method comparison** page quantifies such differences (Chapter 6).

References

- Ahmed, F., & Kilic, K. (2019). Fuzzy analytic hierarchy process: A performance analysis of various algorithms. *Fuzzy Sets and Systems*, 362, 110–128. <https://doi.org/10.1016/j.fss.2018.08.009>
- Beresneva, E., & Gordenko, M. (2018). Applying the methods of system analysis to teaching assistants evaluation. *Proceedings of the Institute for System Programming of the RAS*, 30(3), 251–270. [https://doi.org/10.15514/ISPRAS-2018-30\(3\)-18](https://doi.org/10.15514/ISPRAS-2018-30(3)-18)
- Buckley, J. J. (1985). Fuzzy hierarchical analysis. *Fuzzy Sets and Systems*, 17(3), 233–247. [https://doi.org/10.1016/0165-0114\(85\)90090-9](https://doi.org/10.1016/0165-0114(85)90090-9)
- Kahraman, C., Öztayşi, B., Uçal Sarı, İ., & Turanoğlu, E. (2014). Fuzzy analytic hierarchy process with interval type-2 fuzzy sets. *Knowledge-Based Systems*, 59, 48–57. <https://doi.org/10.1016/j.knosys.2014.02.001>
- Milošević, M. R., Milošević, D. M., Stević, D. M., & Stanojević, A. D. (2019). Smart city: Modeling key indicators in serbia using it2fs. *Sustainability*, 11(13), 3536. <https://doi.org/10.3390/su11133536>
- Saaty, T. L. (1980). *The analytic hierarchy process: Planning, priority setting, resource allocation*. McGraw-Hill.