

LESSON 3 · 15 OCTOBER 2026

Applied Systems Sciences 1

Models, agent-based models and LLMs

How could we find out what saves the lake?

David Maier

Department of Environmental Systems Sciences

University of Graz

Yesterday we fished a lake, and it collapsed

In one or two sentences: **why do fish stocks collapse**, even when no fisher wants them to?

Why lakes collapse

Use what we found: the system, the feedback, the emergence.

How could we find out what saves the lake?

We can't replay the lake, and we can't think it through.

- 1 **How could we find out?** A model.
- 2 **What kind of model?** Agent-based.
- 3 **What if the fishers were LLMs?** And what can an LLM do besides fishing?
- 4 **Building it ourselves.** The tools, and the clinic.

How could we find out what saves the lake?

Hardin (1968): shared resources are doomed, the tragedy of the commons. **Ostrom (1990):** many communities have managed theirs for centuries. Under which conditions?

We can't replay last week's lakes a thousand times, with other rules and other people. And we can't think it through:

“For the complex systems we usually deal with in science, just thinking is not sufficient to rigorously deduce the consequences of our simplifying assumptions: we have to let the computer show us what happens.”

Railsback & Grimm (2019)

Railsback & Grimm (2019), Ch. 1.4 · Biggs et al. (2022), Ch. 1

A model: a purposeful representation

“A model is a purposeful representation of some real system.” Starfield et al. (1990), in Railsback & Grimm (2019)

The **purpose** is the filter: it decides what goes in and what stays out.

“All decisions are made on the basis of models. Most models are in our heads. [...] Mental models have serious shortcomings.” Forrester (1994)

Railsback & Grimm (2019), Ch. 1.2

Same lake, two purposes

In pairs, 4 minutes.

A How many fish can a lake give for ever?


The largest catch per round that the lake can sustain.

B Would talking, or a cooperative, save the lake?

Whether fishers who can agree on rules catch less and earn more.

For each purpose: name one thing you would **include** and one you would **leave out**.

The modelling cycle



Modelling is **iterative**. The first version is always wrong, and that's fine. “It is the very purpose of the model to teach us what is important.”

What kind of model? Two ways to model our lake

An equation

The stock and the total catch. Every fisher is the average fisher.

```
left = stock - catch  
stock = min(2 * left, capacity)
```

Enough for purpose A. It is exactly the rule our game used.

An agent-based model

Every fisher is an agent: own needs, sees what others do, adapts.

```
for fisher in fishers:  
    fisher.fish(lake)
```

Needed for purpose B: trust, talking and differences between fishers.

The price of an ABM: every detail is a choice, and too many details make a model hard to analyse and reproduce.

Railsback & Grimm (2019), Ch. 1.4 · Schlüter et al. (2022), Ch. 28, pp. 385 to 388

What makes a model agent-based?

Agents are **unique** and **autonomous**, and usually **interact locally** with each other and their environment.

1 Unique

Fishers differ: debts, boats, how much they trust others.

2 Autonomous

Each decides their own catch, for their own goals.

3 Local interaction

They talk to their neighbours, not to everyone.

4 Adaptive

They change their catch when the lake or the others change.

Real research on purpose B

Why do cooperatives fail in small-scale fisheries?

Micro

Fishers who differ in reliability; they stay loyal or sell to someone else.

Meso

Fish buyers pick their crew; cooperatives dismiss members and dissolve.

Macro

The fish stock, harvested by all of them together.

Result: initial trust and the diversity among fishers decide which form of organisation establishes itself and lasts.

Lindkvist, Basurto & Schlüter (2017) · Schlüter et al. (2022), Ch. 28, case study 28.1

Quick check

Which of these questions about a fishery needs an agent-based model?

1. How many fish the lake can sustain
2. Whether fishers who trust each other form a cooperative
3. How big last year's total catch was
4. How fast a single fish grows

What if the fishers were LLMs?

Researchers let LLM agents fish a shared lake, month after month, and talk to each other.

“All but the most powerful LLM agents fail to achieve a sustainable equilibrium [...], with the highest survival rate below 54%.” Piatti et al. (2024)

Talking to each other was critical for the lakes that survived.

Piatti et al. (2024), GovSim · Park et al. (2023), generative agents

Role 1: an LLM decides for an agent. The rule is written in language, not code.

```
catch = 3 if lake.stock > 100 else 1
```

“You are a fisher at a shared lake with 80 fish left. How many do you catch this month, and why?”

Why LLM fishers are hard to predict

What an LLM is trained to do

Predict the next word of a text.

Billions of simple units, trained on huge amounts of text.

What it can do

Summarise, translate, write code, play a fisher.

None of these abilities was programmed one by one.

An LLM is itself a system: simple units, a simple objective, abilities nobody wrote down. We study it like our lake, by observing and testing.

Who grouped your answers?

Why lakes collapse

Role 2: an LLM reads what people say. It read your explanations and sorted them into themes. SES research does this with interviews and documents.

What if we group again?

If the LLM groups the same answers again, what will happen?

1. Exactly the same groups
2. Similar groups with different names
3. Completely different groups

Let's group again

Why lakes collapse

Two roles, and the same four questions

An LLM deciding for an agent, or reading what people say, is a model like any other.

Purpose

What question should it answer, and is it the right tool for it?

Assumptions

What does the prompt take for granted? What did the model learn from its training texts?

Validation

Does it match what real fishers did or said?

Reproducibility

Would a second run, or someone else, get the same result?

The same four questions apply to our lake equation and to every model in this course.

Build the lake yourself

Python with uv

Runs the lake model; uv installs Python and packages.

VS Code

Where you write the model, with git and Copilot.

git

Saves every version of your model.

GitHub

Shares the model with your team; you hand in here.

Copilot

AI help in VS Code, free with GitHub Education.

A Groq API key

Lets LLM fishers play, from session 8.

Now: setup clinic. We fix whatever doesn't work yet.

Python and uv

We use **Python 3.12+**, installed and managed by **uv**.

Install uv

Windows (PowerShell):

```
powershell -ExecutionPolicy Bypass -c "irm https://astral.sh/uv/install.ps1 | iex"
```

macOS and Linux:

```
curl -Lsf https://astral.sh/uv/install.sh | sh
```

The two commands you need

```
uv add package-name    # install a package
uv run script.py        # run a script
```

Check

```
uv --version
```

 in a terminal.

VS Code and git

Visual Studio Code

Download: code.visualstudio.com

Windows: the User Installer (64-bit). macOS: the Universal build.

Then install the **Python extension** by Microsoft. VS Code suggests it when you open a .py file.

git

Windows: git-scm.com, the default settings are fine.

macOS: already installed, or installed on first use.

Check: `git --version` in a terminal.

Your Groq API key

For the LLM exercises: free, with limits on requests per minute and per day

1 Sign up

At console.groq.com, on the free plan. No credit card needed.

2 Create a key

Under **API Keys**. Copy it right away.

3 Store it safely

In a `.env` file on your computer. The setup check tests that it works.

A key is like a password. Never paste it into your code, and never commit it to git. Anyone with your key can use your account.

GitHub and Copilot

1 GitHub account

Sign up at github.com with your university e-mail address.

2 GitHub Education

Apply for the Student Developer Pack with your university e-mail. It unlocks Copilot for free.

3 Copilot in VS Code

Sign in to VS Code with your GitHub account. Code suggestions, explanations and fixes.

In the clinic and tonight: follow the setup guide on Moodle, create your Groq key and run the setup check. Post problems in the Moodle forum.

One rule, and how to test it

Take your rule from yesterday that could save the lake. **How would you test it with a model?**

How I would test my rule with a model

One or two sentences. It stays anonymous.

Setup clinic

Work on your own laptop. Ask when something doesn't work.

- 1 **uv** installed: `uv --version` prints a version.
- 2 **VS Code** opens a folder and runs a Python file.
- 3 **git** installed: `git --version` prints a version.
- 4 **Groq API key** created and saved in a `.env` file.
- 5 **GitHub** account, with the Education application sent.

How could we find out what saves the lake?

- 1 We can't replay reality or think it through, so we build a **model**, a purposeful representation.
- 2 Fishers differ and adapt, so an **agent-based model** fits.
- 3 **LLMs** can play the fishers or read what they say, and they are systems we have to test too.

Next week: Thursday 22 October · 08:30 to 09:45 · SR 35.K1 · working with git and GitHub

Sources for this lesson

Railsback, S. F. & Grimm, V. (2019). *Agent-Based and Individual-Based Modeling: A Practical Introduction*. 2nd ed. Ch. 1: models, the modelling cycle, ABM, the rabies example.

Biggs, R. et al. (eds., 2022). *The Routledge Handbook of Research Methods for Social-Ecological Systems*. Schlüter, Lindkvist, Wijermans & Polhill, Ch. 28: agent-based modelling, case study 28.1.

Piatti, G. et al. (2024). Cooperate or collapse: emergence of sustainable cooperation in a society of LLM agents. *NeurIPS 2024*.

Park, J. S. et al. (2023). Generative agents: interactive simulacra of human behavior. *UIST '23*.
doi:10.1145/3586183.3606763

Lindkvist, E., Basurto, X. & Schlüter, M. (2017). Micro-level explanations for emergent patterns of self-governance arrangements in small-scale fisheries. *PLoS ONE* 12(4).

Hardin (1968) · Ostrom (1990) · Forrester (1994)