

GitHub Copilot for data scientists

Less boilerplate.
More time for scientific judgment.

Complete code. Ask better questions.
Delegate bounded work, then verify.

A practical briefing for Desjardins teams

ONE WORKSPACE

01

Explore

Understand data and assumptions

02

Experiment

Build and compare reproducible
baselines

03

Validate

Test code, metrics and leakage
risks

Three ways to work with Copilot

Choose the interaction by how much work you want to delegate.

01 / COMPLETION

Keep typing

Inline suggestions

Finish pandas transforms, SQL queries and test assertions.

You steer every edit.

```
# aggregate transactions  
# by customer and month
```

02 / ASK

Think it through

Conversational help

Explain a notebook, compare metrics or challenge a split.

Advice first; apply deliberately.

```
"Could this feature leak  
future information?"
```

03 / AGENTIC

Delegate a task

Agent + tools

Edit related files, run tests and iterate on a bounded task.

Review diffs and tool actions.

```
"Build a baseline pipeline  
and run its unit tests."
```

More autonomy does not mean less accountability. Keep context scoped and permissions intentional.

RPI: turn a question into verified work

HVE Core on GitHub: Research -> Plan -> Implement, with review and follow-up.

R Research

Understand before editing

Inspect data contracts and code.
Check leakage and alternatives.
Separate evidence from assumptions.

Output: findings + open decisions

`/rpi-research`

P Plan

Make success testable

Choose the split and metrics.
Sequence small, bounded tasks.
Define tests and acceptance criteria.

Output: plan + critique

`/rpi-plan`

I Implement

Execute the approved scope

Build the pipeline and tests.
Run checks; record real results.
Review diffs against the plan.

Output: changes + validation

`/rpi-implement`

Reuse adequate research | Review the plan | Validate results | Route gaps back to the right phase

Right-size the model to the uncertainty

Start with routing, or choose capability manually. Escalate when evidence says you should.

AUTOMATIC ROUTING

Auto / Hydrافusion*

Let the available router
balance
capability, latency and cost.

Use approved options only.
Check the picker and admin policy.

MANUAL CAPABILITY CHOICE

LARGE	Research + rubber ducking Challenge assumptions, compare approaches, look for leakage and counterexamples.
MEDIUM	Planning Turn evidence into tasks, dependencies and measurable acceptance criteria.
SMALL / MEDIUM	Implementation Small for routine edits; medium for complex changes. Escalate if failures or ambiguity persist.

*Hydrافusion is setup-dependent, not a universal Copilot option. Tiers above are a practical heuristic, not product labels.

Try it: a reproducible credit-risk baseline

One synthetic dataset. One clear objective. Three bounded prompts.

01

RESEARCH / ASK

"Inspect this synthetic dataset and notebook.
Identify leakage risks and compare validation splits.
Do not edit files yet."

02

PLAN / HVE RPI

"Plan a reproducible baseline with a time-based split.
Fit preprocessing on training data only.
Define metrics and tests before coding."

03

IMPLEMENT / AGENT

"Implement the approved plan and run the tests.
Report PR-AUC, calibration and subgroup results.
Show the diff and any remaining limitations."

DONE MEANS

A clean rerun works.
Split and leakage tests pass.
Metrics include limitations.
A human reviews the result.

Faster work. Same rigor.

Banking guardrail: synthetic or approved de-identified data only; no client PII or secrets in prompts.