DARTMOUTH ADVISORY PARTNERS

TECHNICAL SHEET / PRIVATE AI

sh4pe

Keep the model close to the work.

TYPE	IN THE MACHINE	PUBLISHED
Local model studio	Checks	October 2026

01 / WHAT IT IS

What it is

sh4pe is a local model studio that trains, tests and runs private AI models on hardware your institution controls. Your knowledge stays inside your walls.

TYPE	Local model studio
GROUP	Private AI
IN THE MACHINE	Checks
SOLVES	Private AI and model evaluation, Institutional intelligence
BROUGHT BY	Advisory

02 / HOW IT WORKS

How it works

In order. Each step is checked before the next begins.

01	**Bounded dataset**
02	**Local training**
03	**Held-out evaluation**
04	**Qualified use**

> IN PRACTICE
>
> A candidate model is assessed on a defined task before it can take a place in a governed workflow. Local deployment alone does not establish that authority.

03 / WHO IT IS FOR, AND HOW IT IS APPLIED

Who it is for, and how it is applied

Who it is for

Teams working with sensitive institutional knowledge and explicit data boundaries.

How it is applied

Model quality and the permitted task are evaluated separately from the choice to run locally. A model takes on a task only after it passes a test on data it has never seen, and a person signs off.

04 / THE FIGURES

The figures

The result, on a sealed exam

MEASURE	RESULT
The open-weight model, untrained	34% (70 of 203)
The same model, after training in sh4pe	76% (154 of 203, strict match)
Training time, on one machine	5 minutes (300 seconds)
Exam	sipp-holdout200-k, 203 questions the model never saw
Answers, by grade	151 exact, 3 normalized, 21 partial, 28 misses
Seal check	0 of 203 exam clauses appear in the training set

Measured 25 September 2026 on six public Canadian pension policy documents. Qwen2.5 3B, exam sipp-holdout200-k. The full report is eval_report.json in the adapter folder.

The training recipe, as recorded

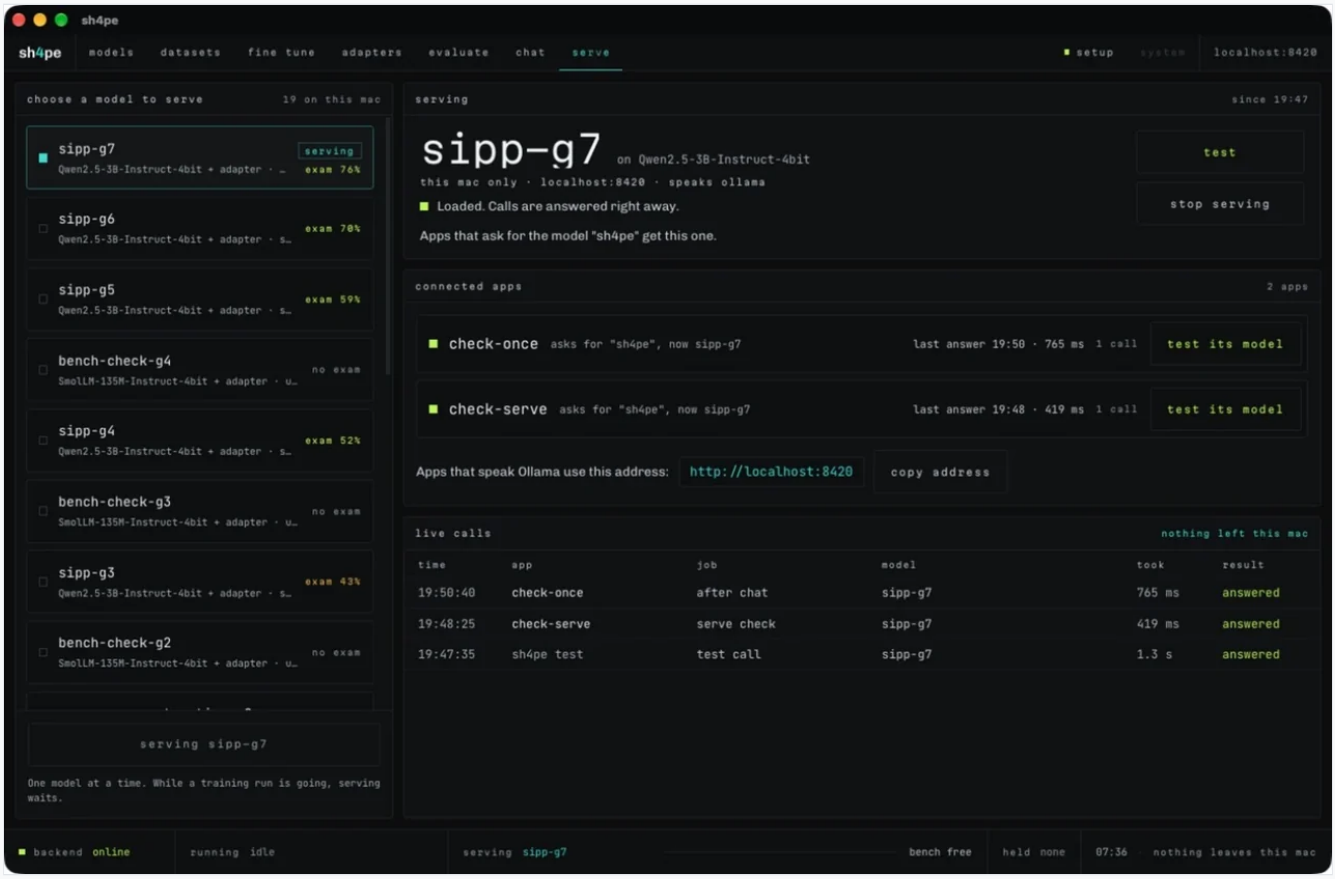
SETTING	VALUE
Base model	mlx-community/Qwen2.5-3B-Instruct-4bit
Adapter	sipp-g7 (LoRA rank 8)
Dataset	sipp-extraction-v6
Iterations	1,050
Batch size	1
Learning rate	5.00e-5
Trainable parameters	6.652 million
Checkpoints	10, every 100 iterations, 25.4 MB each, kept beside the final weights

Source: the sipp-g7 session record in sh4pe, trained 25 September 2026.

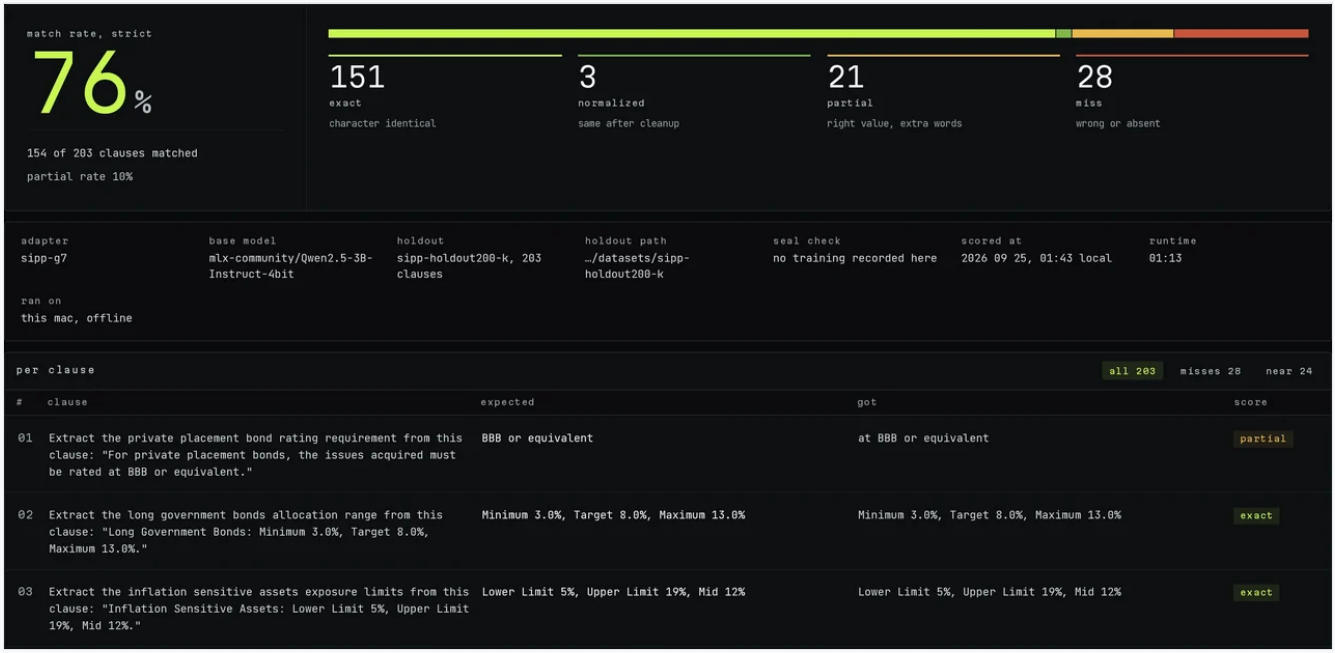
05 / THE SCREENS

The screens

Real captures of the working app. Nothing is mocked up.



sh4pe serving a trained specialist to other apps on the same machine. Every call is logged. Its text never is.



match rate, strict

76%

154 of 203 clauses matched

partial rate 10%

151	3	21	28
exact	normalized	partial	miss
character identical	same after cleanup	right value, extra words	wrong or absent

adapter	base model	holdout	holdout path	seal check	scored at	runtime
sipp-g7	mlx-community/Qwen2.5-3B-Instruct-4bit	sipp-holdout200-k, 203 clauses	…/datasets/sipp-holdout200-k	no training recorded here	2026 09 25, 01:43 local	01:13

ran on: this mac, offline

per clause — all 203 | misses 28 | near 24

#	clause	expected	got	score
01	Extract the private placement bond rating requirement from this clause: "For private placement bonds, the issues acquired must be rated at BBB or equivalent."	BBB or equivalent	at BBB or equivalent	partial
02	Extract the long government bonds allocation range from this clause: "Long Government Bonds: Minimum 3.0%, Target 8.0%, Maximum 13.0%."	Minimum 3.0%, Target 8.0%, Maximum 13.0%	Minimum 3.0%, Target 8.0%, Maximum 13.0%	exact
03	Extract the inflation sensitive assets exposure limits from this clause: "Inflation Sensitive Assets: Lower Limit 5%, Upper Limit 19%, Mid 12%."	Lower Limit 5%, Upper Limit 19%, Mid 12%	Lower Limit 5%, Upper Limit 19%, Mid 12%	exact

The scorecard: 76% strict match on 203 clauses, each clause's key and answer side by side.



session record · sipp-g7 — before and after | close

adapter	sipp-g7	session	ses_8e8c2d896d10	outcome	completed
base model	mlx-community--Qwen2.5-3B-Instruct-4bit	dataset	sipp-extraction-v6	dataset path	datasets/sipp-extraction-v6
iterations	1050	learning rate	5.00e-5	lora rank	8
batch size	1	trainable params	6.652 m	trained	2026 09 25, 01:42
anchors captured	2 before and after				

measured

76% strict match on sipp-holdout200-k, 203 clauses. Exact 151, normalized 3, partial 21, miss 28.

Scored 2026 09 25, 01:43 local, 01:13 on this machine. The full report is written to eval_report.json in the adapter folder.

exams, every report in the folder

sipp-holdout200-k — 76% exact 151, normalized 3, partial 21, miss 28 of 203 — 2026 09 25, 01:43

checkpoints on disk

100 · 25.4 MB | 200 · 25.4 MB | 300 · 25.4 MB | 400 · 25.4 MB | 500 · 25.4 MB | 600 · 25.4 MB | 700 · 25.4 MB | 800 · 25.4 MB | 900 · 25.4 MB | 1000 · 25.4 MB

Saved by mlx-lm during the run and kept beside the final weights.

loss curve, as recorded — train — val — train 0.151 at 1050 — val 1.07 at 1050

One training run, read from its own folder: recipe, score, checkpoints and loss curve.

06 / CONNECTED WORK

Connected work

sh4pe works alongside these systems.

SYSTEM	WHAT IT DOES
Antechamber	Test the authority before granting it.
Overture	Authority before action.

07 / READ ON

Read on

1. Living Policy Architecture, on SSRN
   `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=7108198`
2. The same paper, read on dap.solutions
   `https://dap.solutions/papers/living-policy-architecture/`
3. sh4pe on dap.solutions
   `https://dap.solutions/solutions/systems/sh4pe/`

Govern first. Then automate.

Models propose. Humans ratify. To discuss sh4pe, write to info@dap.solutions.

DARTMOUTH ADVISORY PARTNERS / TORONTO, CANADA / DAP.SOLUTIONS



DARTMOUTH

ADVISORY PARTNERS

Your organization's foundations, compiled.