

Qanat

Experiment Tracking for HPC Users

Ammar Mian

November 11, 2025

1. Introduction

2. Core Workflow

3. Power Features

4. Demo Examples

5. Wrap-up

Computational research faces a critical challenge:

- “Which parameters gave that result 3 months ago?”

Computational research faces a critical challenge:

- “Which parameters gave that result 3 months ago?”
- “What was the exact code version for that experiment?”

Computational research faces a critical challenge:

- “Which parameters gave that result 3 months ago?”
- “What was the exact code version for that experiment?”
- “How do I reproduce this on the cluster?”

Computational research faces a critical challenge:

- “Which parameters gave that result 3 months ago?”
- “What was the exact code version for that experiment?”
- “How do I reproduce this on the cluster?”

The Challenge

Existing tools (MLflow, Weights & Biases) require GUIs and databases, but HPC users work in **headless terminal environments** with job submission systems.

What HPC Users Need:

- CLI-first workflow
- Works on remote servers
- HTCondor/Slurm integration
- Minimal dependencies
- Offline operation

What HPC Users Need:

- CLI-first workflow
- Works on remote servers
- HTCondor/Slurm integration
- Minimal dependencies
- Offline operation

What GUI Tools Require:

- Web browser access
- Database servers
- Network connectivity
- Complex setup
- Heavy dependencies

What HPC Users Need:

- CLI-first workflow
- Works on remote servers
- HTCondor/Slurm integration
- Minimal dependencies
- Offline operation

What GUI Tools Require:

- Web browser access
- Database servers
- Network connectivity
- Complex setup
- Heavy dependencies

Qanat's Solution

A minimal, CLI-based experiment tracking system designed specifically for computational research on HPC systems.

```

1  -----
2  |   _   |                               |   _   |
3  | | | | | _ _ _ _ _ |   _   | | | |
4  | | | | | / _ ' | ' _ \ / _ ' |   _   |
5  \ \ \ ' / ( _ | | | | | ( _ | | | |
6  \ _ / \ _ \ \ _ _ , _ | | | | \ _ _ , _ | \ _ _ |
7
8  Qanat, version 1.9.4
9  Minimal experience management system.

```

- A tool I built for my own research needs
- Used by some students, but nothing groundbreaking
- Everything it does *can* be done without it
- Power users could write it in 50 lines of bash

So why use it? It enforces a structured workflow and handles the boring parts automatically.

1. Introduction

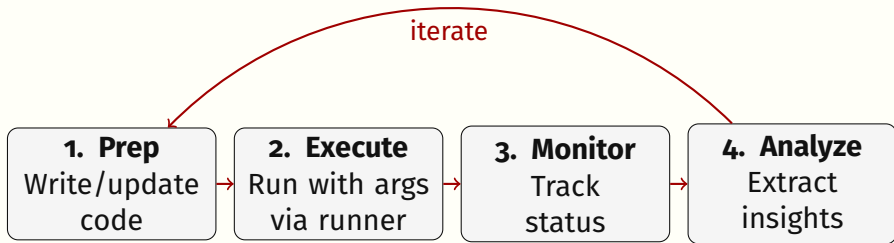
2. Core Workflow

3. Power Features


4. Demo Examples

5. Wrap-up

Standard steps in any computational experiment:



Qanat organizes these steps and forces you to work in a reproducible way.



Automatic tracking: git commit, parameters, stdout/stderr, duration, status

Let's see qanat in action!

We'll demonstrate:

1. Initialize a project
2. Create a simple training script
3. Register an experiment
4. Run with parameters
5. Explore results interactively

Demo Note

Everything shown works the same on your laptop or on a 1000-node cluster.

For every run, qanat stores:

```
1 results/my_experiment/run_42/  
2 |-- info.yaml           # All parameters + metadata  
3 |-- stdout.txt          # Complete output  
4 |-- stderr.txt          # Errors & warnings  
5 '-- [your results]      # Whatever your script saves
```

info.yaml contains:

- Exact git commit SHA
- All command-line arguments
- Execution date, duration, status
- Tags, description, dataset links

1. Introduction

2. Core Workflow

3. **Power Features**

4. Demo Examples

5. Wrap-up

The Problem: Running experiments over multiple parameter combinations

The Problem: Running experiments over multiple parameter combinations

Qanat's Solution: Groups

The Problem: Running experiments over multiple parameter combinations

Qanat's Solution: Groups

- -g flag: Define parameter groups

The Problem: Running experiments over multiple parameter combinations

Qanat's Solution: Groups

- -g flag: Define parameter groups
- -r flag: Define ranges with start/stop/step

The Problem: Running experiments over multiple parameter combinations

Qanat's Solution: Groups

- -g flag: Define parameter groups
- -r flag: Define ranges with start/stop/step
- Cartesian product between groups and ranges

Hyperparameter search across learning rates and seeds:

```
1 qanat experiment run train_model \  
2   --epochs 100 \  
3   -g "--lr 0.001" \  
4   -g "--lr 0.01" \  
5   -g "--lr 0.1" \  
6   -r "--seed 0 5 1"
```

Hyperparameter search across learning rates and seeds:

```
1 qanat experiment run train_model \  
2   --epochs 100 \  
3   -g "--lr 0.001" \  
4   -g "--lr 0.01" \  
5   -g "--lr 0.1" \  
6   -r "--seed 0 5 1"
```

This creates a **single run_id** with 15 groups (3 LRs × 5 seeds)

Hyperparameter search across learning rates and seeds:

```
1 qanat experiment run train_model \  
2   --epochs 100 \  
3   -g "--lr 0.001" \  
4   -g "--lr 0.01" \  
5   -g "--lr 0.1" \  
6   -r "--seed 0 5 1"
```

This creates a **single run_id** with 15 groups (3 LRs × 5 seeds)

```
1 results/train_model/run_17/  
2 |-- group_0/  # lr=0.001, seed=0  
3 |-- group_1/  # lr=0.001, seed=1  
4 |-- ...  
5 '-- group_14/ # lr=0.1, seed=4
```


Benefits of the groups concept:

- **Aggregation:** K-fold cross-validation, multiple random seeds
- **Organization:** Related experiments stay together
- **Efficiency:** Submit all variations in one command
- **Analysis:** Actions can operate on entire groups

Design Choice

Groups let you think about parameter *sweeps* rather than individual runs.

Runner Abstraction: Same command, different execution environments

Runner	Use Case	Backend
local	Laptop/workstation	subprocess
htcondor	HTCondor clusters	Python bindings
slurm	Slurm clusters	sbatch

Just add: `-runner htcondor`

Everything else stays the same!

HTCondor Note

Requires `htcondor` Python bindings <11 (version 10.x recommended)

Submit to cluster:

```
1 qanat experiment run train_model \  
2   --runner htcondor \  
3   --submit_template gpu_template \  
4   -g "--lr 0.001" -g "--lr 0.01" -g "--lr 0.1" \  
5   -r "--seed 0 5 1"
```

Monitor progress:

```
1 qanat experiment status train_model --live
```

Qanat creates submit descriptions, tracks job status, and aggregates results automatically.

Reusable resource specifications:

Define once, use everywhere:

- CPU cores, memory requirements
- GPU requests (number, type)
- Docker/Singularity images
- Queue names, time limits

Example

`gpu_template`: Requests 1 GPU, 16GB RAM, cuda11.8 environment

Switch between small CPU tests and large GPU runs with just the template name.

Reproducible environments with Singularity/Apptainer:

- `-container path/to/image.sif`
- Automatic bind mounting of working directory
- Dataset paths mounted automatically
- GPU support via `-nv` flag
- Works with both local and HPC runners

Reproducibility++

Container + git commit = exact code *and* exact environment.

```
1 # Run in container on cluster with GPU
2 qanat experiment run deep_learning \
3   --runner htcondor \
4   --container /shared/images/pytorch.sif \
5   --gpu True \
6   -g "--model resnet50" \
7   -g "--model efficientnet" \
8   -r "--batch_size 16 64 16"
```

Qanat handles:

- Bind mounting code and data
- GPU passthrough (-nv)
- Submit description generation
- Tracking container path in metadata

The most important feature: Qanat refuses to run uncommitted code!

Enforced Reproducibility

If your code isn't committed, your experiment isn't reproducible. Period.

Every run automatically tracks:

- Exact git commit SHA (e.g., a3f4b2c)
- Branch name
- Repository remote URL
- Working directory state

This is the foundation that makes everything else meaningful.

Qanat handles git checkout automatically!

Scenario 1: Exact reproduction

```
1 # Reruns with same commit + same parameters
2 qanat experiment rerun my_experiment 42
```

Scenario 2: Old code, new parameters

```
1 # Check which commit was used
2 qanat experiment run_explore my_experiment
3 # Shows: git_commit: a3f4b2c
4
5 # Run with that commit + different parameters
6 qanat experiment run my_experiment \
7   --commit_sha a3f4b2c \
8   --new_param value
```


Reproducing Experiments]Reproducing Experiments
No manual git checkout needed! Qanat does it for you.

Problem: Don't want to rerun experiments just to change a plot.

Problem: Don't want to rerun experiments just to change a plot.

Solution: Actions - Post-processing scripts that analyze stored results

Problem: Don't want to rerun experiments just to change a plot.

Solution: Actions - Post-processing scripts that analyze stored results

- Receive `-storage_path` automatically from qanat
- Access all files saved during the run
- Generate plots, compute metrics, export data
- Can be run from CLI or interactive menu

Problem: Don't want to rerun experiments just to change a plot.

Solution: Actions - Post-processing scripts that analyze stored results

- Receive `-storage_path` automatically from qanat
- Access all files saved during the run
- Generate plots, compute metrics, export data
- Can be run from CLI or interactive menu

Benefit: Separate computation from visualization. Run experiment once, analyze many ways.

Define action in experiment YAML:

```
1 actions:
2   - plot_curves:
3       executable: scripts/plot_learning_curves.py
4       executable_command: python
5       description: Plot training/validation curves
```

Define action in experiment YAML:

```
1 actions:
2   - plot_curves:
3       executable: scripts/plot_learning_curves.py
4       executable_command: python
5       description: Plot training/validation curves
```

Action script structure:

```
1 import argparse, numpy as np, matplotlib.pyplot as plt
2
3 parser = argparse.ArgumentParser()
4 parser.add_argument("--storage_path", required=True)
5 args = parser.parse_args()
6
7 # Load results from storage_path
8 results = np.load(f"{args.storage_path}/metrics.npz")
9 plt.plot(results['train_loss'], label='train')
10 plt.savefig(f"{args.storage_path}/learning_curves.png")
```

Define action in experiment YAML:

```
1 actions:
2   - plot_curves:
3       executable: scripts/plot_learning_curves.py
4       executable_command: python
5       description: Plot training/validation curves
```

Action script structure:

```
1 import argparse, numpy as np, matplotlib.pyplot as plt
2
3 parser = argparse.ArgumentParser()
4 parser.add_argument("--storage_path", required=True)
5 args = parser.parse_args()
6
7 # Load results from storage_path
8 results = np.load(f"{args.storage_path}/metrics.npz")
9 plt.plot(results['train_loss'], label='train')
10 plt.savefig(f"{args.storage_path}/learning_curves.png")
```

Execute: qanat experiment action my_exp plot_curves 42

Real-time experiment monitoring:

- **TQDM integration:** Training loops show progress bars
- **Count format:** Progress: 42/100 for parallel jobs
- **Group aggregation:** See progress across all parameter combinations
- **Live dashboard:** `-live` flag for auto-refreshing status

Implementation

Your script writes progress to stdout in recognized formats. Qanat parses and displays it.

1. Introduction

2. Core Workflow

3. Power Features

4. Demo Examples

5. Wrap-up

Simple C program with parallel processing:

```
1  # Run with different thread counts
2  qanat experiment run montecarlo_pi --n_trials 10000000 \
3    -g "--num_threads 1" -g "--num_threads 2" -g "--num_threads 4" -g "
    --num_threads 8"
```

Simple C program with parallel processing:

```
1 # Run with different thread counts
2 qanat experiment run montecarlo_pi --n_trials 10000000 \
3   -g "--num_threads 1" -g "--num_threads 2" -g "--num_threads 4" -g "
   --num_threads 8"
```

Qanat tracks:

- Progress: Each thread writes count to `progress.txt`
- Results: Pi estimate, error, execution time
- Parameters: `n_trials`, `num_threads`

Simple C program with parallel processing:

```
1 # Run with different thread counts
2 qanat experiment run montecarlo_pi --n_trials 10000000 \
3   -g "--num_threads 1" -g "--num_threads 2" -g "--num_threads 4" -g "
   --num_threads 8"
```

Qanat tracks:

- Progress: Each thread writes count to `progress.txt`
- Results: Pi estimate, error, execution time
- Parameters: `n_trials`, `num_threads`

Plot action (gnuplot):

```
1 qanat experiment action montecarlo_pi plot 42
```

Creates convergence plot + random sampling visualization

PyTorch training with TensorBoard:

```
1 # Compare CNN vs pretrained EfficientNet
2 # Each experiment has its own dedicated script
3 uv run qanat experiment run mnist_cnn \
4     --epochs 10 --lr 0.001 \
5     -r "--seed 0 3 1" # 3 random seeds
6
7 uv run qanat experiment run mnist_efficientnet \
8     --epochs 5 --lr 0.0001 \
9     -r "--seed 0 3 1"
```

PyTorch training with TensorBoard:

```
1 # Compare CNN vs pretrained EfficientNet
2 # Each experiment has its own dedicated script
3 uv run qanat experiment run mnist_cnn \
4     --epochs 10 --lr 0.001 \
5     -r "--seed 0 3 1" # 3 random seeds
6
7 uv run qanat experiment run mnist_efficientnet \
8     --epochs 5 --lr 0.0001 \
9     -r "--seed 0 3 1"
```

Automatic tracking:

- Progress: Updates after each epoch (count format)
- TensorBoard logs: Saved to storage_path/tensorboard/
- Metrics: Train/val/test accuracy and loss

PyTorch training with TensorBoard:

```
1 # Compare CNN vs pretrained EfficientNet
2 # Each experiment has its own dedicated script
3 uv run qanat experiment run mnist_cnn \
4     --epochs 10 --lr 0.001 \
5     -r "--seed 0 3 1" # 3 random seeds
6
7 uv run qanat experiment run mnist_efficientnet \
8     --epochs 5 --lr 0.0001 \
9     -r "--seed 0 3 1"
```

Automatic tracking:

- Progress: Updates after each epoch (count format)
- TensorBoard logs: Saved to storage_path/tensorboard/
- Metrics: Train/val/test accuracy and loss

If you can do this in 50 lines of bash, why use qanat?

If you can do this in 50 lines of bash, why use qanat?

- **Enforces discipline:** Forces git commits, no shortcuts
- **Handles boilerplate:** Job submission, directory creation, logging
- **Consistency:** Everyone on your team uses the same structure
- **Lower barrier:** Students don't need to be bash experts
- **Documentation:** Forgot how something works? Check old runs

If you can do this in 50 lines of bash, why use qanat?

- **Enforces discipline:** Forces git commits, no shortcuts
- **Handles boilerplate:** Job submission, directory creation, logging
- **Consistency:** Everyone on your team uses the same structure
- **Lower barrier:** Students don't need to be bash experts
- **Documentation:** Forgot how something works? Check old runs

Reality

It's just organized file management and script wrappers. Nothing revolutionary, but it saves time.

1. Introduction
2. Core Workflow
3. Power Features
4. Demo Examples
5. **Wrap-up**

What qanat is:

- A personal tool I made to stop losing my own experiments
- Used by a few students in my group
- Open source on GitHub (AGPLv3)
- Maintained when I need new features

What qanat is NOT:

- Not polished or well-documented
- Not feature-complete (Slurm support is half-baked)
- Not tested beyond my own use cases
- Not a replacement for proper experiment management systems

Reality

It works for what I need. Might work for you. Might not. Your mileage may vary.

Known limitations:

- Slurm support still in development
- Interactive menu requires Linux/macOS (no Windows support)
- SQLite database limits concurrent writes
- No built-in visualization dashboard

Potential improvements:

- Better dataset versioning
- Experiment comparison tools
- Integration with Jupyter notebooks
- Enhanced action system

1. **Git commit tracking is essential:** Without it, nothing else matters

1. **Git commit tracking is essential:** Without it, nothing else matters
2. **Forced discipline saves you later:** No uncommitted code = no regrets

1. **Git commit tracking is essential:** Without it, nothing else matters
2. **Forced discipline saves you later:** No uncommitted code = no regrets
3. **Checkout old commits to reproduce:** Time-travel debugging actually works

1. **Git commit tracking is essential:** Without it, nothing else matters
2. **Forced discipline saves you later:** No uncommitted code = no regrets
3. **Checkout old commits to reproduce:** Time-travel debugging actually works
4. **HPC job submission is just boilerplate:** Qanat handles the boring parts

1. **Git commit tracking is essential:** Without it, nothing else matters
2. **Forced discipline saves you later:** No uncommitted code = no regrets
3. **Checkout old commits to reproduce:** Time-travel debugging actually works
4. **HPC job submission is just boilerplate:** Qanat handles the boring parts
5. **Nothing revolutionary:** Just organized wrappers, but that's often enough

1. **Git commit tracking is essential:** Without it, nothing else matters
2. **Forced discipline saves you later:** No uncommitted code = no regrets
3. **Checkout old commits to reproduce:** Time-travel debugging actually works
4. **HPC job submission is just boilerplate:** Qanat handles the boring parts
5. **Nothing revolutionary:** Just organized wrappers, but that's often enough

Bottom Line

If you're already disciplined with git and bash scripts, you don't need this. If not, it might help.

Documentation & Code:

- **Documentation:** <https://ammarmian.github.io/qanat/>
- **GitHub:** <https://github.com/ammarmian/qanat>
- **Installation:** `pip install git+https://github.com/ammarmian/qanat`

Questions?

Feedback and contributions welcome!

Thank You!

Happy experimenting on your HPC clusters

`ammam.mian@univ-smb.fr`