



Seqera AI: How we're using LLMs and Agents with Nextflow code

Creating systems that help scientists to build and run complex workflows



Phil Ewels
Senior Product Manager for OSS
phil.ewels@seqera.io

Outline

- 01.** Introduction
- 02.** Seqera AI Chat
- 03.** Tool integration
- 04.** AI agents
- 05.** Moving fast
- 06.** Looking ahead



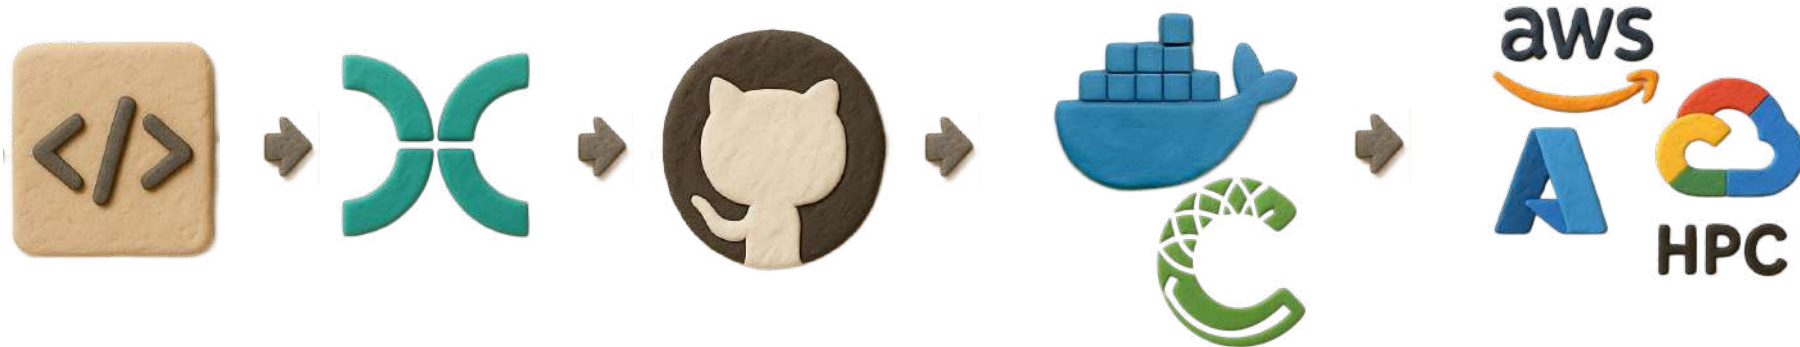
Introduction

01



Nextflow

Open-source orchestrator for writing & deploying workflows at scale.



Nextflow

Open-source orchestrator for writing & deploying workflows at scale.

- Programming language for bioinformatics workflows
- Orchestrator for executing workflows
- Runs on everything, with data from anywhere
- Reproducibility built-in, native support for software packaging and code version control



Scalable

Portable

Reproducible

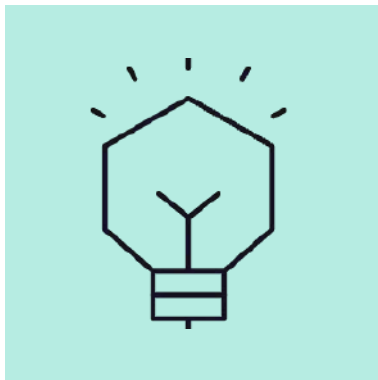
Seqera AI Chat

02



What's the problem?

What do Nextflow users complain about the most?



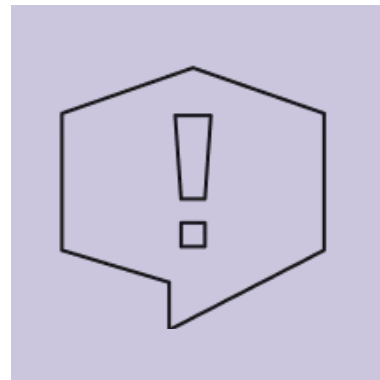
Steep learning curve

It's Groovy, baby



Syntax changed between versions

DSL1 → DSL2



Difficult / unusual errors

Generalists need not apply

Base models struggle with Nextflow

- Nextflow questions are a small subset of data science
- Answers are often very specific to Nextflow and not generally applicable
- Can't retrain the base model (too expensive)
- Can't include a long enough prompt (context window)

My Nextflow pipeline exited with an error in the FastQC process, saying exit code 127. What should I do?


Exit code 127 usually means a command not found error.

Check if it's installed:

```
fastqc --version
```

If it's not found, you need to install it, for example using conda:

```
conda install -c bioconda fastqc
```

Please don't tell people to do that.. 

Base models struggle with writing Nextflow*

- Most public Nextflow code is old
- DSL1 → DSL2 syntax change in 2022+
- Hallucinates syntax from other (more common) languages
- Isn't as opinionated as we'd like it to be!

* Though they're getting better all the time!



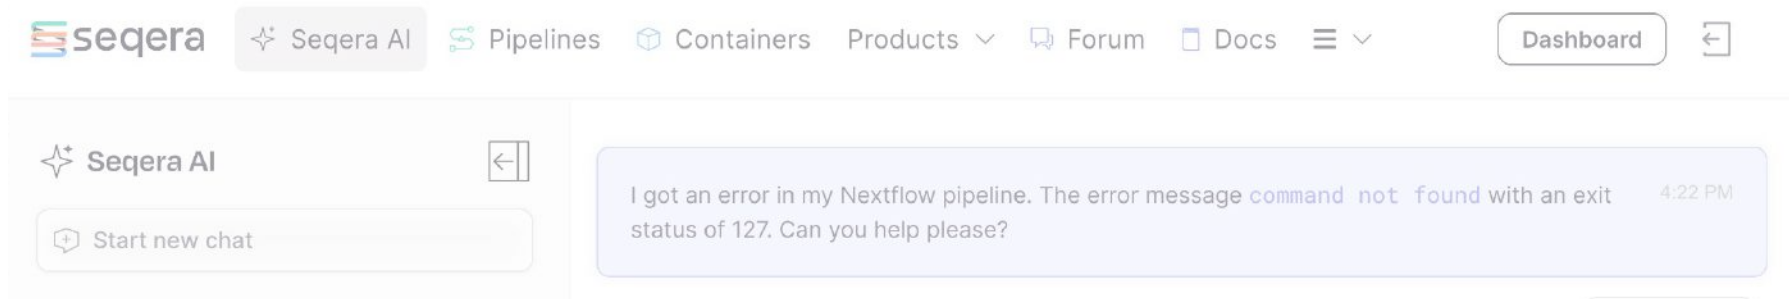
Seqera AI

Bioinformatics AI that helps you get from 0 to 1 for all your omics

<https://seqera.io/ask-ai/>

- Latest Nextflow documentation is prioritised over everything else
- Also includes docs from nf-core, nf-test and other trusted sources*
- Prioritises nf-core best practices and knows available pipelines

* Not Slack



The screenshot displays the Seqera AI web interface. At the top, there is a navigation bar with the Seqera logo and links for Seqera AI, Pipelines, Containers, Products, Forum, Docs, and a Dashboard button. Below the navigation bar, the main content area features a chat window titled 'Seqera AI'. The chat window has a 'Start new chat' button and a message from the AI assistant. The message reads: 'I got an error in my Nextflow pipeline. The error message `command not found` with an exit status of 127. Can you help please?'. The message is timestamped '4:22 PM'.



Seqera AI

Bioinformatics AI that helps you get from 0 to 1 for all your omics

<https://seqera.io/ask-ai/>

[demo]



Seqera AI

Pipelines

Containers

Products

Forum

Docs



Dashboard



Seqera AI



Start new chat

I got an error in my Nextflow pipeline. The error message `command not found` with an exit status of 127. Can you help please?

4:22 PM



Tool integration

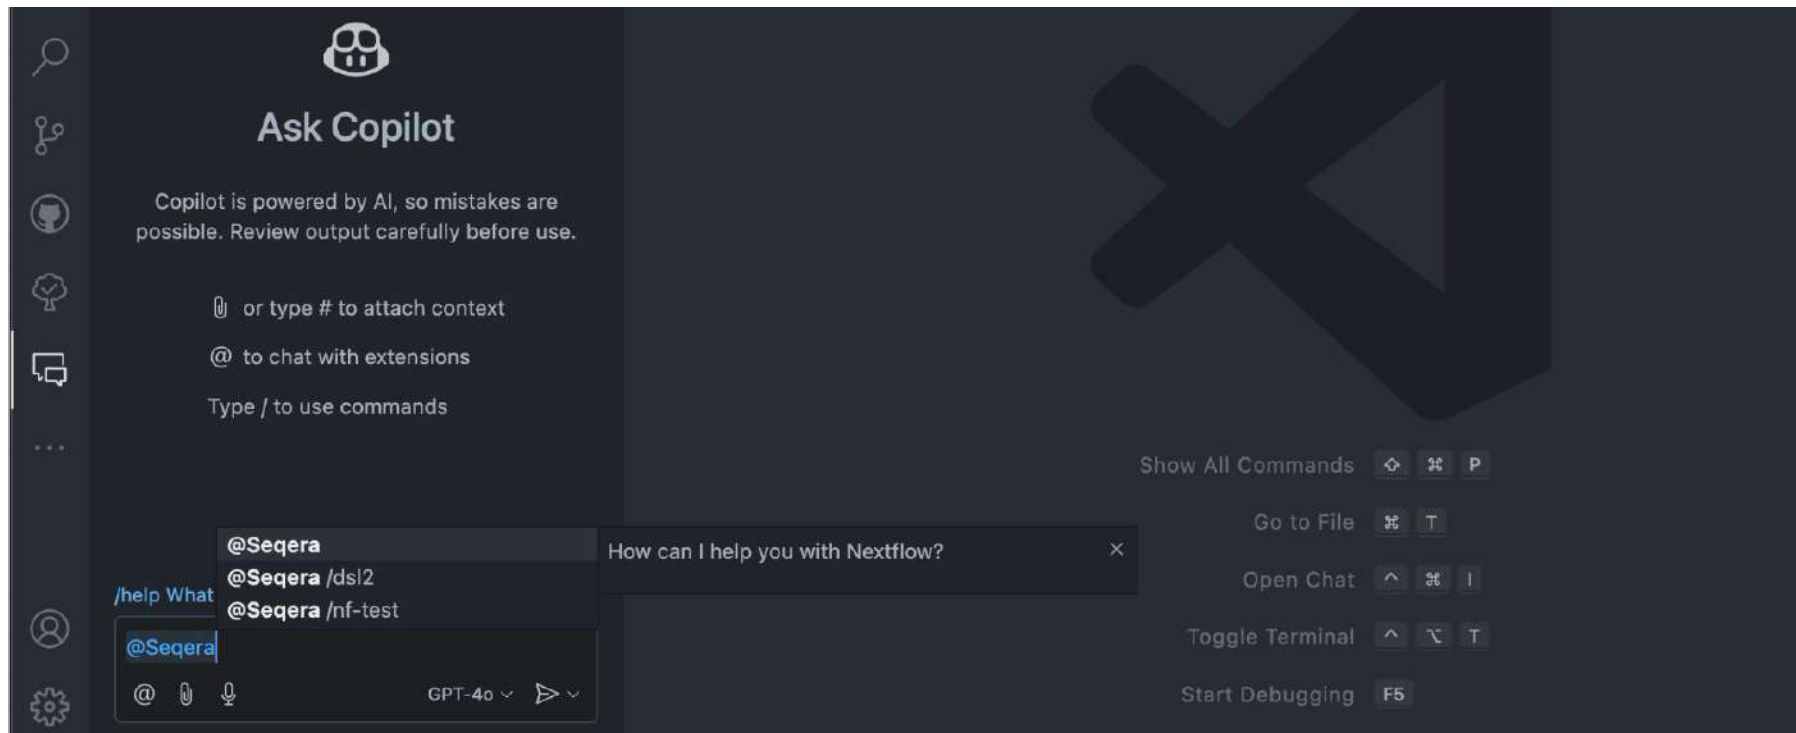
03



Seqera AI in VS Code

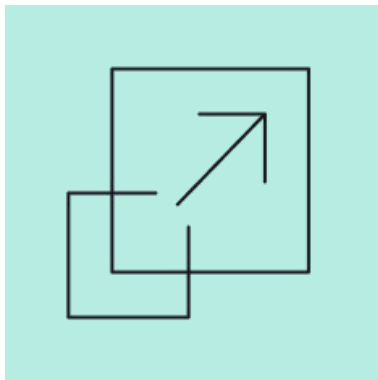
Nextflow help right where you need it

<https://nextflow.io/vscode>



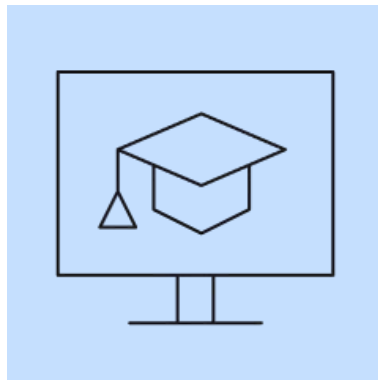
What's the problem?

What's most difficult about MultiQC reports?



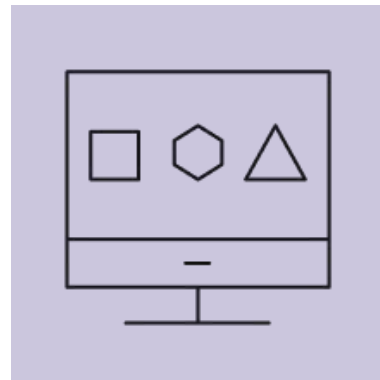
Large sample numbers

Big data is..big



Understanding results

Is this wiggle normal?



Seeing the big picture

Detecting patterns

Seqera AI in MultiQC

Analysis help right where you need it

<https://docs.seqera.io/multiqc/ai>

- Help from AI, in the context of a report
- Generation at report time, or on demand
- Flexible, can be used with most providers

multiqc
v1.27.dev0

General Stats

STAR

Summary Statistics

Alignment Scores

multiqc

A modular tool to aggregate results from bioinformatics analyses across many samples into a single report.

Report generated on 2025-01-21, 23:53 CET based on data in:
/Users/ewels/GitHub/MultiQC/website/public/examples/rna-seq/data

✦ Summarize Report

Toolbox



AI agents

04



The problem with AI code generation

AI interns can create more work than they solve

- AI writes code, but doesn't understand it
- Looks impressive until you try to run it
- Now you're debugging an entire pipeline looking for the error, rather than iteratively writing + testing

```
ERROR ~ No signature of method:
groovyx.gpars.dataflow.DataflowBroadcast.into() is
applicable for argument types:
(Script_85c4d82870d584dc$_runScript_closure1) values:
[Script_85c4d82870d584dc$_runScript_closure1@53aa2fc9]
Possible solutions: any(), find(), bind(java.lang.Object),
any(groovy.lang.Closure), find(groovy.lang.Closure),
is(java.lang.Object)

-- Check script 'demo.nf' at line: 6 or see
'.nextflow.log' file for more details
```



Testing Nextflow code isn't trivial

- Proper modularity and code structure
- Finding example data
- Using the nf-test framework
- AI agents can help



AI agents

Going beyond prompts and responses

Agency Level	Description	What that's called	Example Pattern
☆☆☆	LLM output has no impact on program flow	Simple Processor	<code>process_llm_output(llm_response)</code>
☆☆☆	LLM output determines an if/else switch	Router	<code>if llm_decision(): path_a() else: path_b()</code>
☆☆☆	LLM output determines function execution	Tool Call	<code>run_function(llm_chosen_tool, llm_chosen_args)</code>
☆☆☆	LLM output controls iteration and program continuation	Multi-step Agent	<code>while llm_should_continue(): execute_next_step()</code>
☆☆☆	One agentic workflow can start another agentic workflow	Multi-Agent	<code>if llm_trigger(): execute_agent()</code>

Source: https://huggingface.co/docs/smolagents/en/conceptual_guides/intro_agents



Moving fast

06



Keeping pace with progress

How Seqera AI adapts to a rapidly changing landscape

Claude 4 release timeline (yesterday)

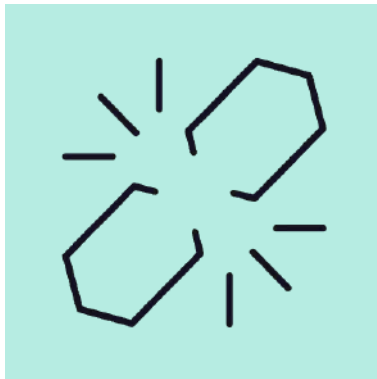
- 7:01 PM - Claude 4 announced on X/Twitter
- 7:16 PM - Started testing internally
- 9:06 PM - Testing finished, good results
- 10:18 PM - Sonnet 4 live in production



Sasha
@Sasha248
Claude 4 live with Sonnet 4 in prod. nice work team;
@PHH - the pipelines demo should be much more snappy
now

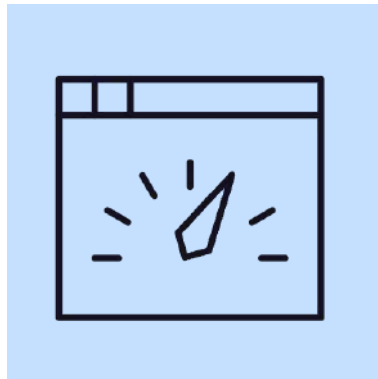
But how can we move so fast?

Updating with confidence



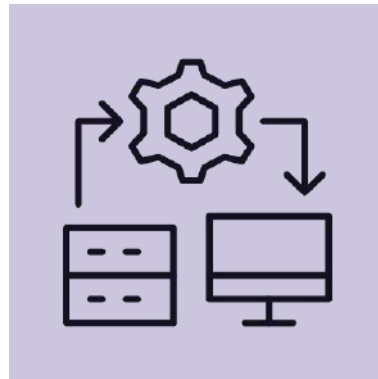
Decouple

Tools like Langchain decouple application logic from LLM providers



Evaluate

Testing rigs give objective metrics about how well a given setup works



Deploy

Robust and automated deployment infrastructure allows fast deployment and rollback (if needed)

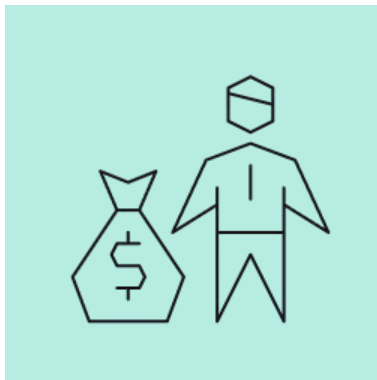
Looking ahead

07



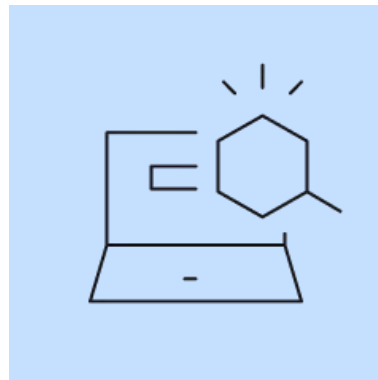
Helpful LLMs

Approaches with Seqera AI



Use the big models

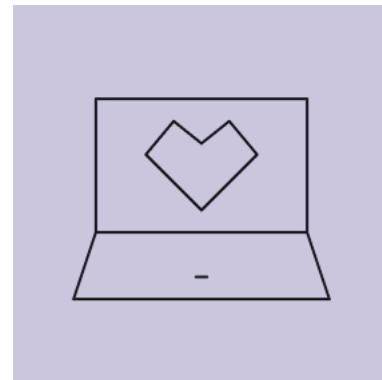
It's impossible to win against the big model providers, better to use them: standardise and be flexible.



RAG and prompts

No-one else can do RAG (retrieval augmented generation) quite like you..

Fine-tuned LLMs can provide additional value over major providers.



Integration is king

Bringing LLMs out of a chat window and into applications helps to make them part of people's workflows.

What's next for Seqera AI

- Configure and launch pipelines
- Help when things go wrong
- Reduce the learning curve
- More integrations for AI agents
- Do all this in a way that is open, trusted and transparent

nf-core/rnaseq:
11 sections with 113 configurable fields

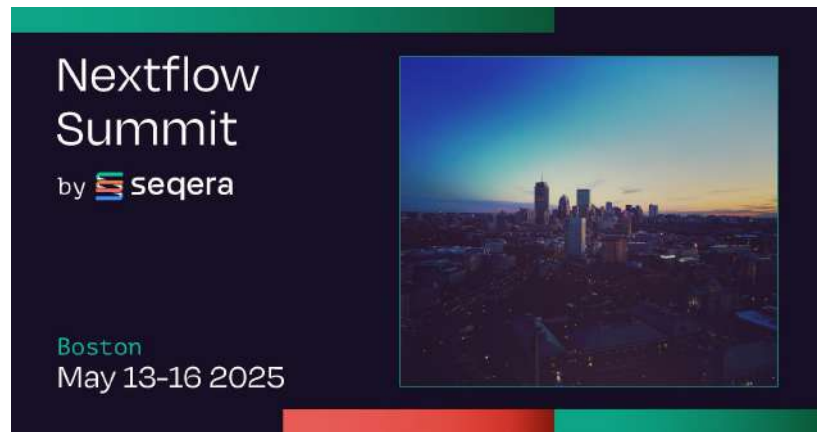
- Input/Output Options: 4
- Reference Genome Options: 21
- Read Trimming Options: 4
- Read Filtering Options: 5
- UMI Options: 9
- Alignment Options: 16
- Optional Outputs: 10
- Quality Control: 6
- Process Skipping Options: 20
- Institutional Config Options: 6
- Generic Options: 12



Find out more



<https://seqera.io/podcasts/>



<https://summit.nextflow.io/2025/boston/>



Thank you

<https://seqera.io/ask-ai/>

<https://docs.seqera.io/multiqc/ai>

<https://nextflow.io/vscode>



nextflow.io



<https://nf-co.re>



seqera.io/multiqc