Symphonic HiGHS

Operationalizing next moves with DecisionOps

Ryan O'Neil · June 27, 2024

👋 Hello, I'm Ryan O'Neil

Currently: Nextmv co-founder & CTO

Building a DecisionOps platform for OR practitioners

Previously: Grubhub, Zoomer, MITRE

Led decision engineering teams, built many a model

Likes cats, cellos, and camping

I also make excellent llama jokes in my spare time

This talk is about

🤠 How Nextmv started using HiGHS and what we use it for

🚀 How HiGHS works in Netmv's DecisionOps platform to provide scalable open-source optimization

🤔 Some meanderings about OSS and optimization

How we HiGHS

How Nextmv has used HiGHS so far

💰 Large-ish custom models for incentive allocation in Go (~4m row x ~6m column MIPs)

📅 Shift scheduling and order fulfillment apps in Go

🐈‍⬛ Price optimization and facility location apps in AMPL

🐾 AMPL partners built stochastic facility location models using AMPL, HiGHS, Nextmv, and Streamlit

HiGHS Shift Scheduling

HiGHS MIP Go Scheduling

Solve a shift scheduling problem with the low-code Nextmv Shift Scheduling app. Define available workers and open shifts, and then run the app to get an assigned shift plan.

Clone app | Customize app | GitHub

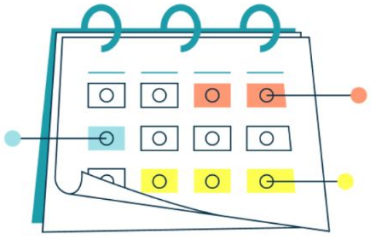

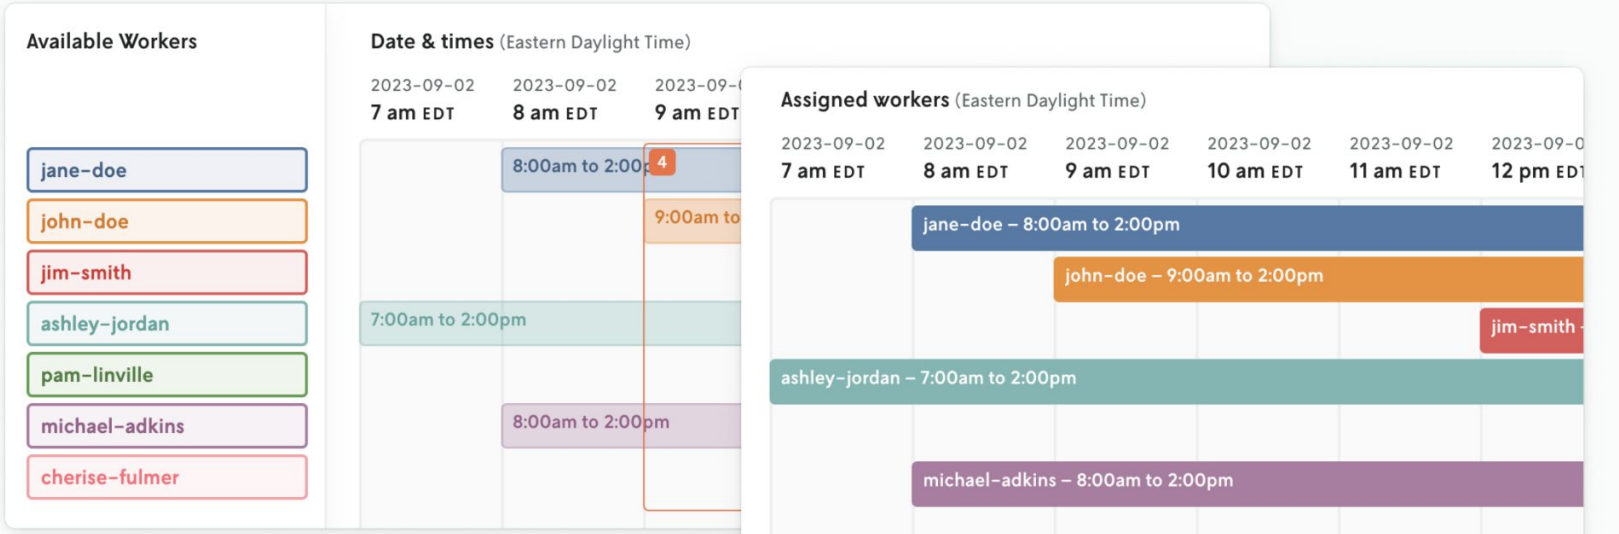

HiGHS Shift Scheduling

Assign workers to shifts, minimizing overages and maximizing objectives.

HiGHS MIP Go Scheduling

HiGHS Order Fulfillment

Solves a MIP-based Order Fulfillment model with the Nextmv SDK.

HiGHS MIP Go Fulfillment

AMPL Price Optimization

Solves a price optimization Mixed Integer Programming problem using AMPL.

HiGHS AMPL Python Pricing

AMPL Facility Location

Solves a facility location problem using AMPL.

HiGHS AMPL Python Facility Location

DecisionOps with 

🔭 Let’s explore DecisionOps with HiGHS and Nextmv

- Leading open source optimization solver
- Applicable to many use cases and problem types

- Platform to build, test, deploy, and operate models efficiently
- Integrates with many optimization solutions

- **Deploy and run custom decision models on prod-ready infrastructure**
- **Perform experiments, share results, manage versions, monitor runs**

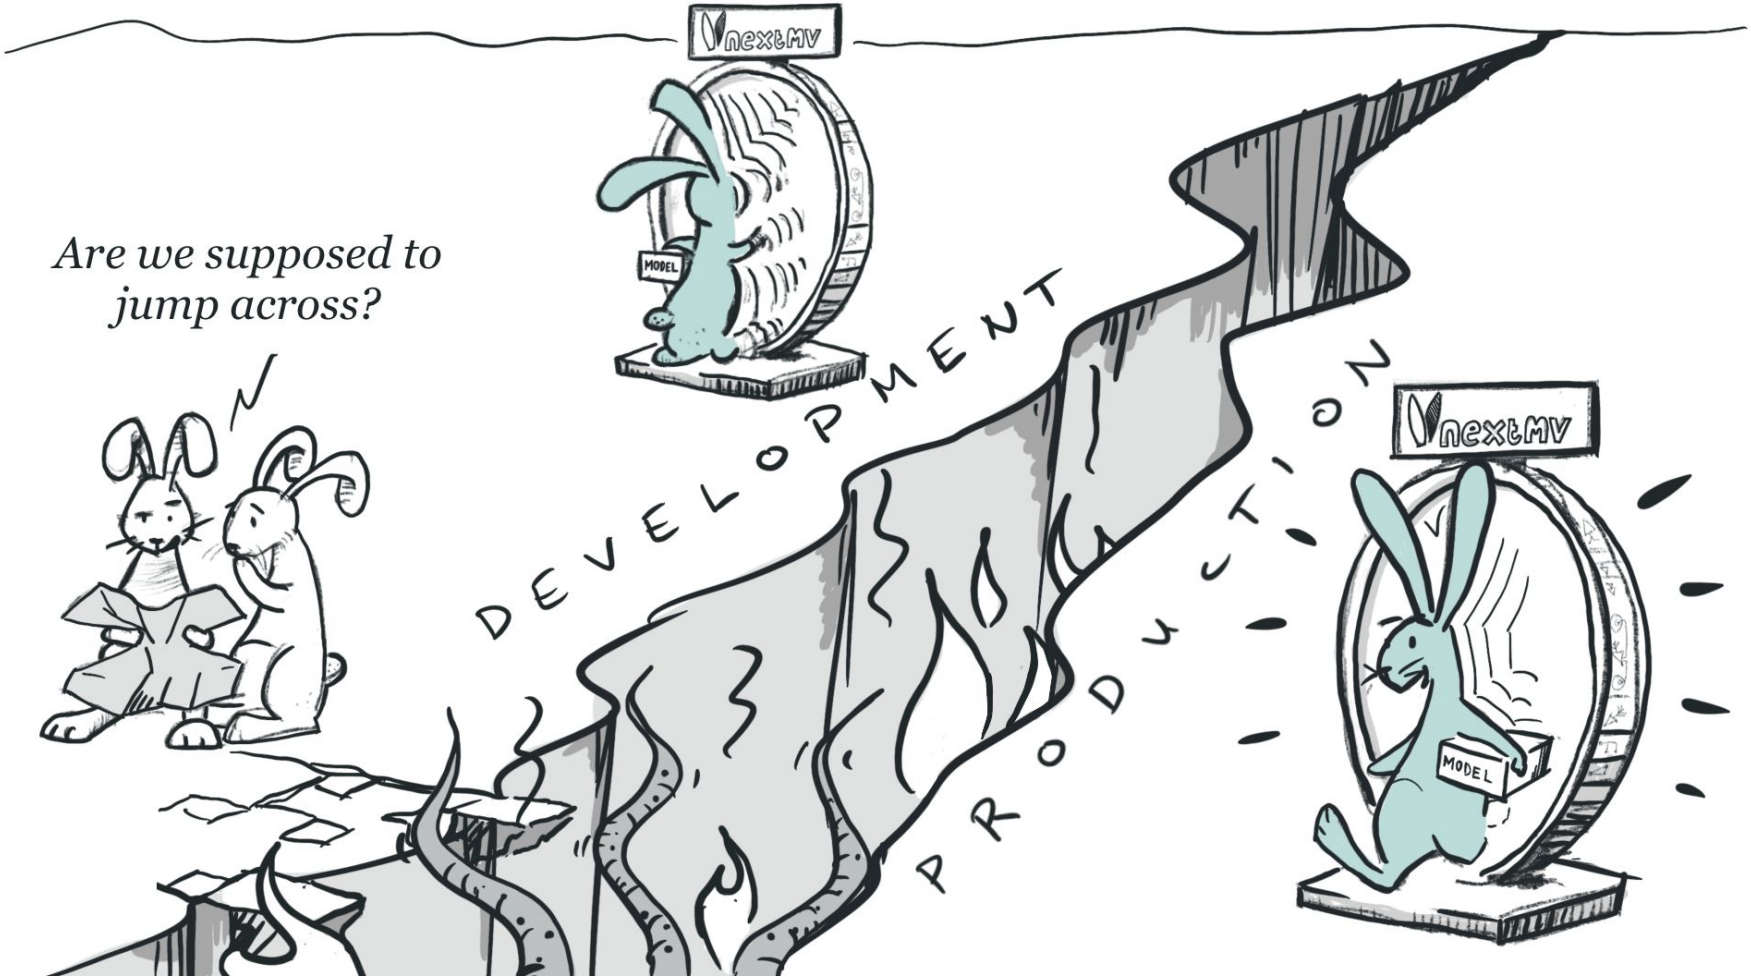
Are we supposed to jump across?
nextmv
MODEL
DEVELOPMENT
PRODUCTION
nextmv
MODEL

🧐 The workflow as we've known it

Elicit requirements from stakeholders

Translate requirements into solver speak

Test decision model in a solver (how?!?)

Hire a developer to wrap it in a microservice

Make that microservice work in the cloud

Live test (what?!? how!?!) 😱

Field questions from unhappy operators

$$\begin{aligned} &\text{maximize} && c^T x \\ &\text{subject to} && Ax \leq b \\ & && x \geq 0 \\ & && x \in Z^n \end{aligned}$$

Takes weeks to months,
involves multiple teams

🧪 Testing MIP model formulations

Traditionally

- ➔ Assemble test data on local machine
- ➔ Write custom code in notebook
- ➔ Set up, perform comparisons
- ➔ Translate to slides for collaboration

Ideally

- ➔ Have usable test sets ready to go
- ➔ Deploy/run model code remotely
- ➔ Tests are repeatable
- ➔ Results are sharable, consistent

A workflow with DecisionOps

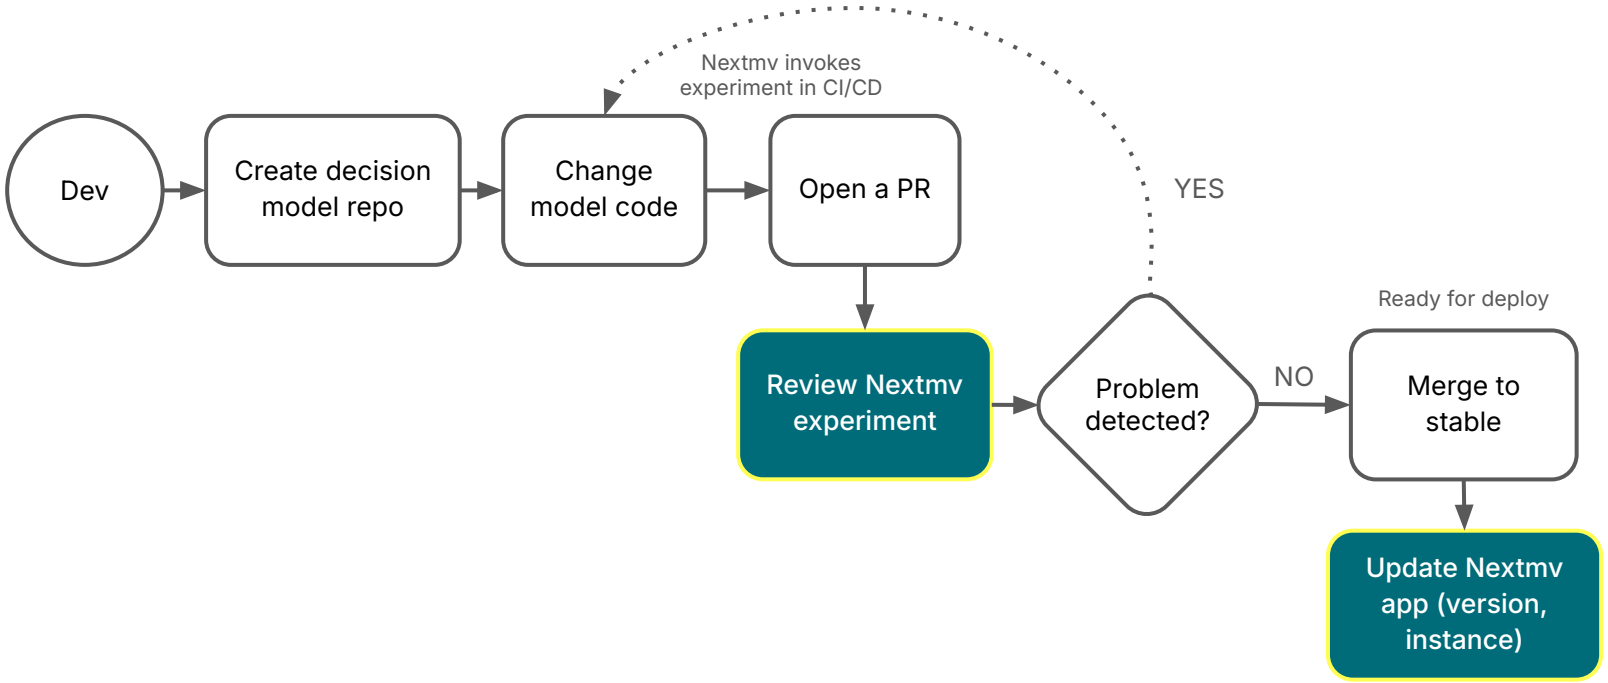

Decision model testing framework

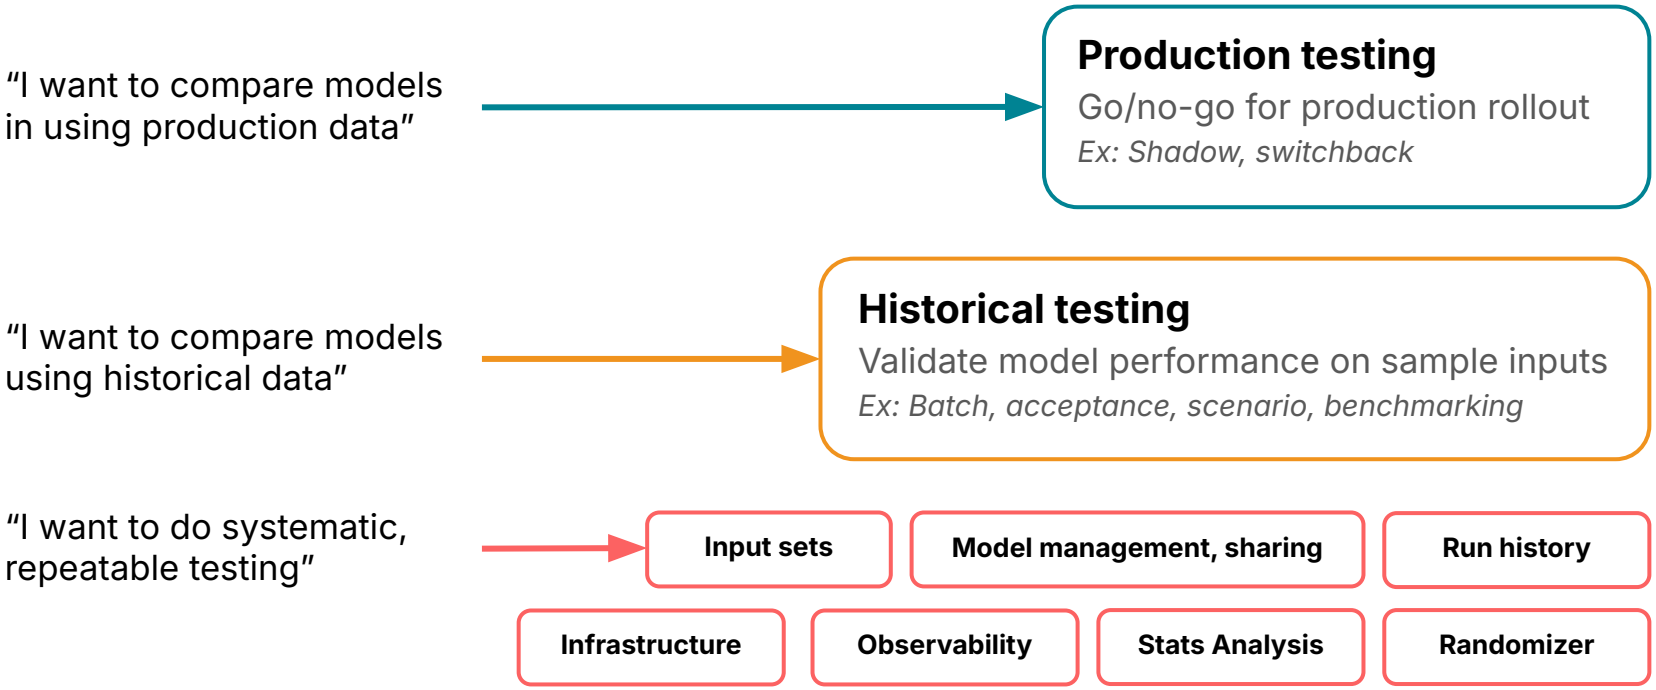

The Nextmv DecisionOps platform

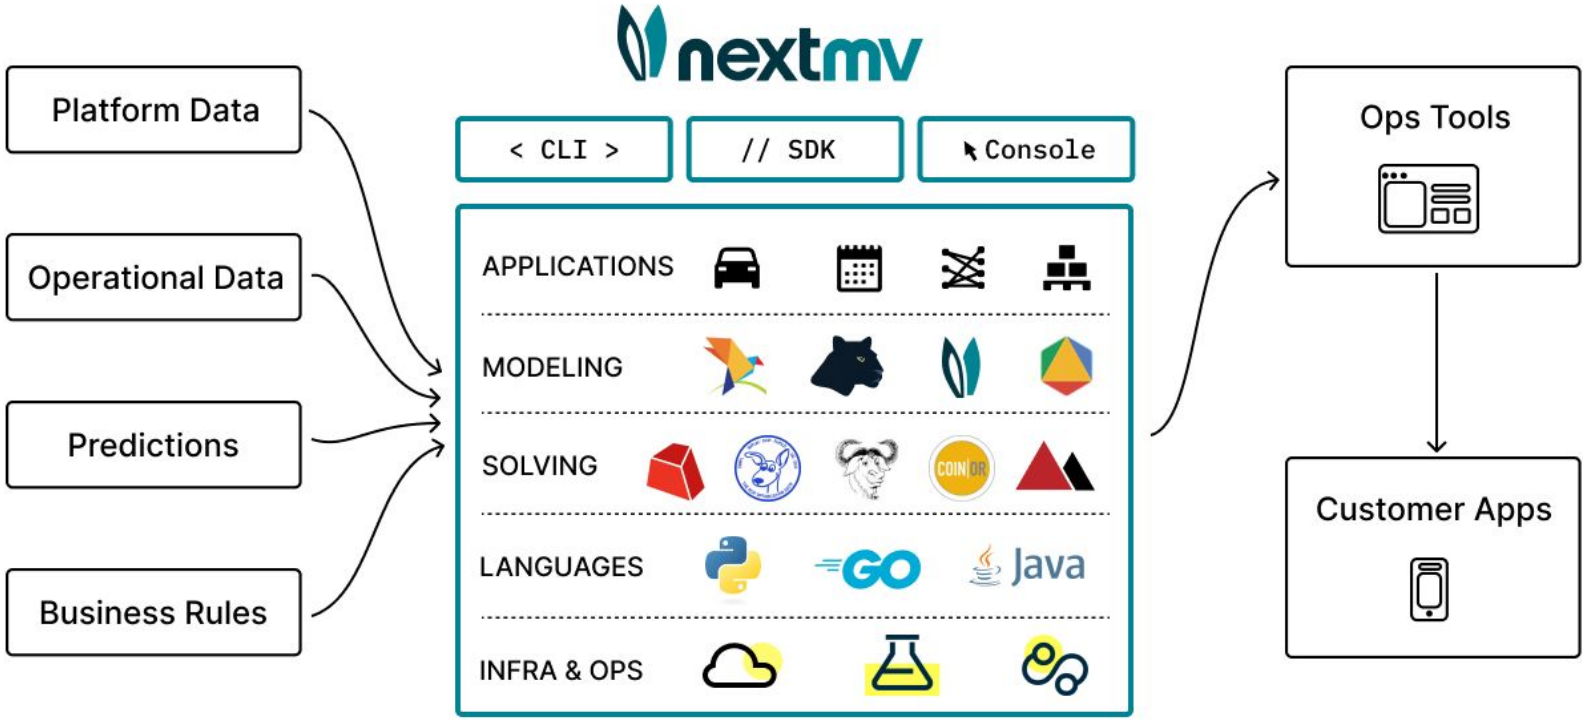

🚜 Today’s example: The Farm Share Company

- Consumer delivery service for farm-based goods
- Recently expanded to offer avocados (yeah, toast!)
- We want to price and supply avocados so as to maximize profits (while considering waste)
- How do we factor uncertainty into our transportation cost estimates?

nextmv speedrun

😮 What just happened?

- We cloned and ran HiGHS-based models for scheduling, order fulfillment, and facility location.
- We fit a price elasticity curve to the avocado data set.
- We ran a scenario test using that curve and different transportation costs in a HiGHS-based pricing model.

Test workflow in practice

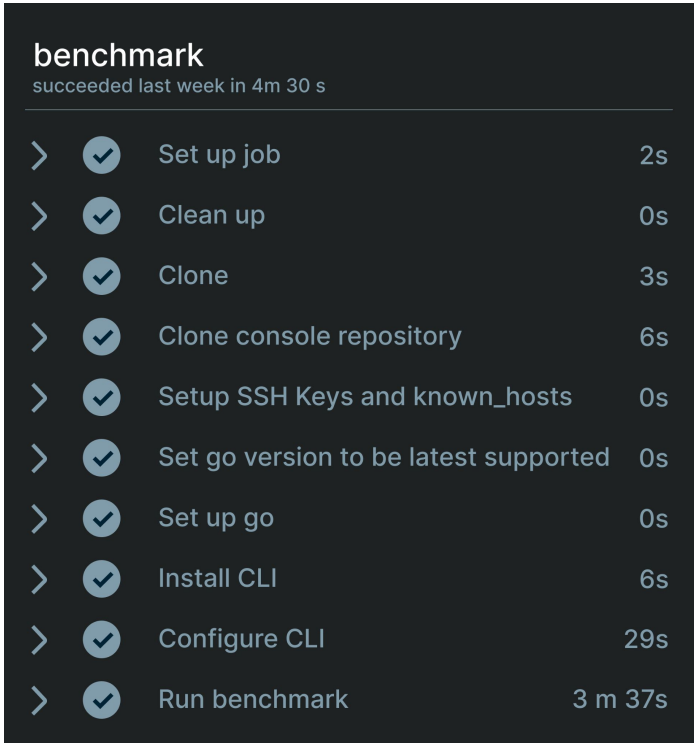

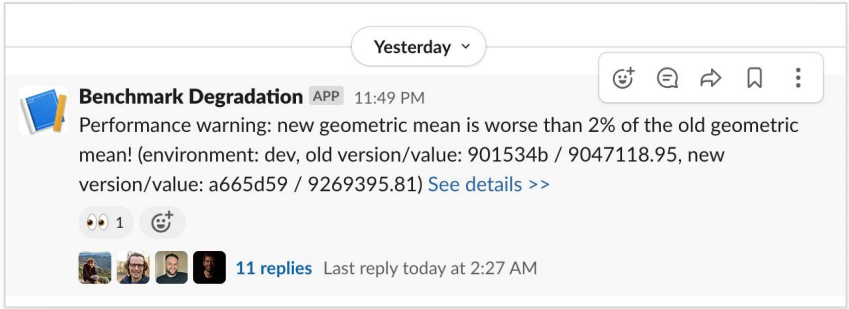

OSS and optimization

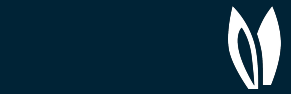

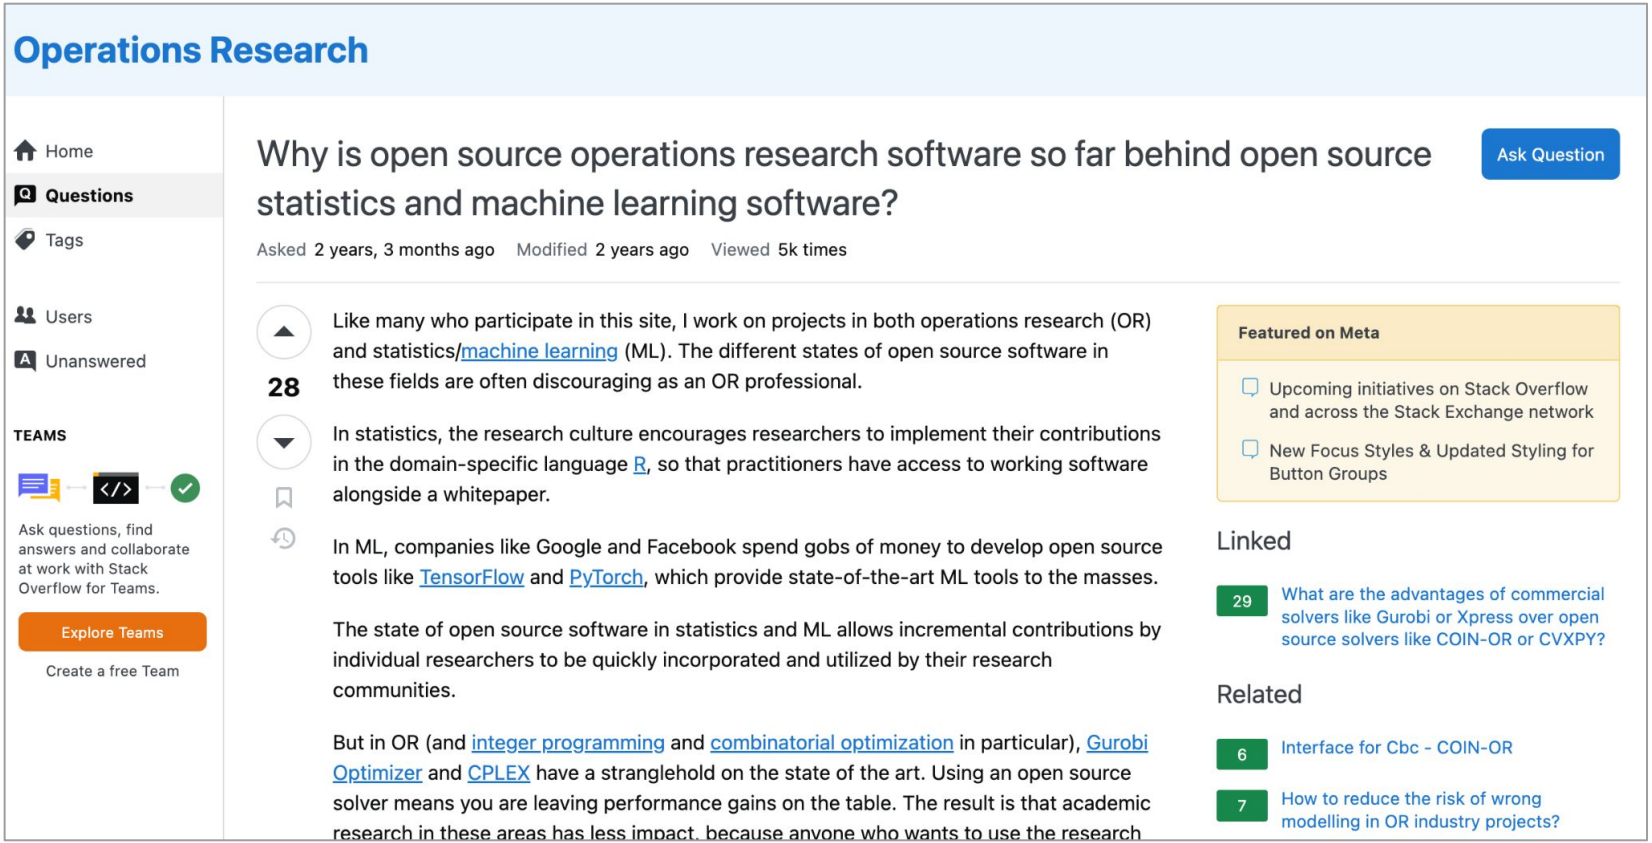
Operations Research
Home
Questions
Tags
Users
Unanswered
TEAMS
Ask questions, find answers and collaborate at work with Stack Overflow for Teams.
Explore Teams
Create a free Team
Why is open source operations research software so far behind open source statistics and machine learning software?
Ask Question
Asked 2 years, 3 months ago Modified 2 years ago Viewed 5k times
28
Like many who participate in this site, I work on projects in both operations research (OR) and statistics/machine learning (ML). The different states of open source software in these fields are often discouraging as an OR professional.
In statistics, the research culture encourages researchers to implement their contributions in the domain-specific language R, so that practitioners have access to working software alongside a whitepaper.
In ML, companies like Google and Facebook spend gobs of money to develop open source tools like TensorFlow and PyTorch, which provide state-of-the-art ML tools to the masses.
The state of open source software in statistics and ML allows incremental contributions by individual researchers to be quickly incorporated and utilized by their research communities.
But in OR (and integer programming and combinatorial optimization in particular), Gurobi Optimizer and CPLEX have a stranglehold on the state of the art. Using an open source solver means you are leaving performance gains on the table. The result is that academic research in these areas has less impact, because anyone who wants to use the research
Featured on Meta
Upcoming initiatives on Stack Overflow and across the Stack Exchange network
New Focus Styles & Updated Styling for Button Groups
Linked
29 What are the advantages of commercial solvers like Gurobi or Xpress over open source solvers like COIN-OR or CVXPY?
Related
6 Interface for Cbc - COIN-OR
7 How to reduce the risk of wrong modelling in OR industry projects?


RepoTracker
1h ago
pytorch
pytorch
0
30d
90d
6m
1y
all
Commits (30d)
Daily Commits (excl. bots)
Latest (excl. bots)
CaoE · 7h ago
Set simdlen based on ATEN_CPU_CAPABILITY (#123514)...
xuhancn · 13h ago
fix miss isa bool check (#128274)
cyyever · 14h ago
[5/N] Remove unused functions. (#127185)
peterbell10 · 16h ago
Don't create python tuple when _maybe_handle_torch_function is called from C++...
Pull Requests (30d)
Daily Opened Pull Requests (excl. bots)
Latest (excl. bots)
OPEN bdhirsh · 1h ago
add TORCH_FORCE_SYNCHRONOUS_COLLECTIVES to force functional collectives to be synchronous
OPEN xuhancn · 1h ago
[Don't merge] Try to restructure code
OPEN jamesjwu · 1h ago
Support aliasvoided modules and op overloads in AOTAutogradCache
OPEN ezyang · 2h ago
Add some guard to size oblivious has_internal_overlap
Contributors (all)
Contributions (excl. bots)
5d EMA
14d EMA
Release
RepoTracker
21m ago
google
or-tools
1
30d
90d
6m
1y
all
Commits (30d)
Daily Commits (excl. bots)
Latest (excl. bots)
ron-at-swgy · 4d ago
Update Python CMake file for OpenBSD (fr4266)
lperron · 5d ago
Merge pull request #4262 from google/dependabot/pip/bazel/pip-9d732ab251...
ron-at-swgy · 5d ago
Use correct fp_env definition for OpenBSD (fr4259)
lperron · 6d ago
Merge pull request #4257 from ron-at-swgy/OpenBSD-Updates...
Pull Requests (30d)
Daily Opened Pull Requests (excl. bots)
Latest (excl. bots)
CLOSED ron-at-swgy · 2d ago
Update SWIG flags for OpenBSD
CLOSED varshneydevansh · 5d ago
replace CapSub with CapAdd for StartMin value
CLOSED ron-at-swgy · 5d ago
Use correct fp_env definition for OpenBSD
CLOSED ron-at-swgy · 6d ago
Updates for OpenBSD build support
Contributors (all)
Contributions (excl. bots)
5d EMA
14d EMA
Release

🦘SCIP 20 Workshop: “SCIP: Past, Present, Future”

Industry has a rather different answer:

“Resoundingly yes!”

Future Needs

TECHNISCHE UNIVERSITÄT DARMSTADT

Does the world need an academic solver like SCIP?

Maybe no:

- Researchers are increasingly using Gurobi (no other solver) for experiments.
- The gap is increasing (see last slide).
- One reason: careful engineering, which is hard to do in a research project (e.g., tuning parallelization).

Well, yes:

- We need a framework, e.g., for branch-and-cut-and-price.
- It is essential to have the source code for developing MIP- or MINLP- solvers.
- It is scientificly sound to at least be able to understand what an underlying solver is doing instead of treating it as a black-box.

🙌 Of course we need high quality open source MIP

- Lots of companies use HiGHS, SCIP, OR-Tools, and others in production.
- Data science and software engineering users are particularly used to adopting and deploying OSS tools.
- We really should use more OSS optimization in research.

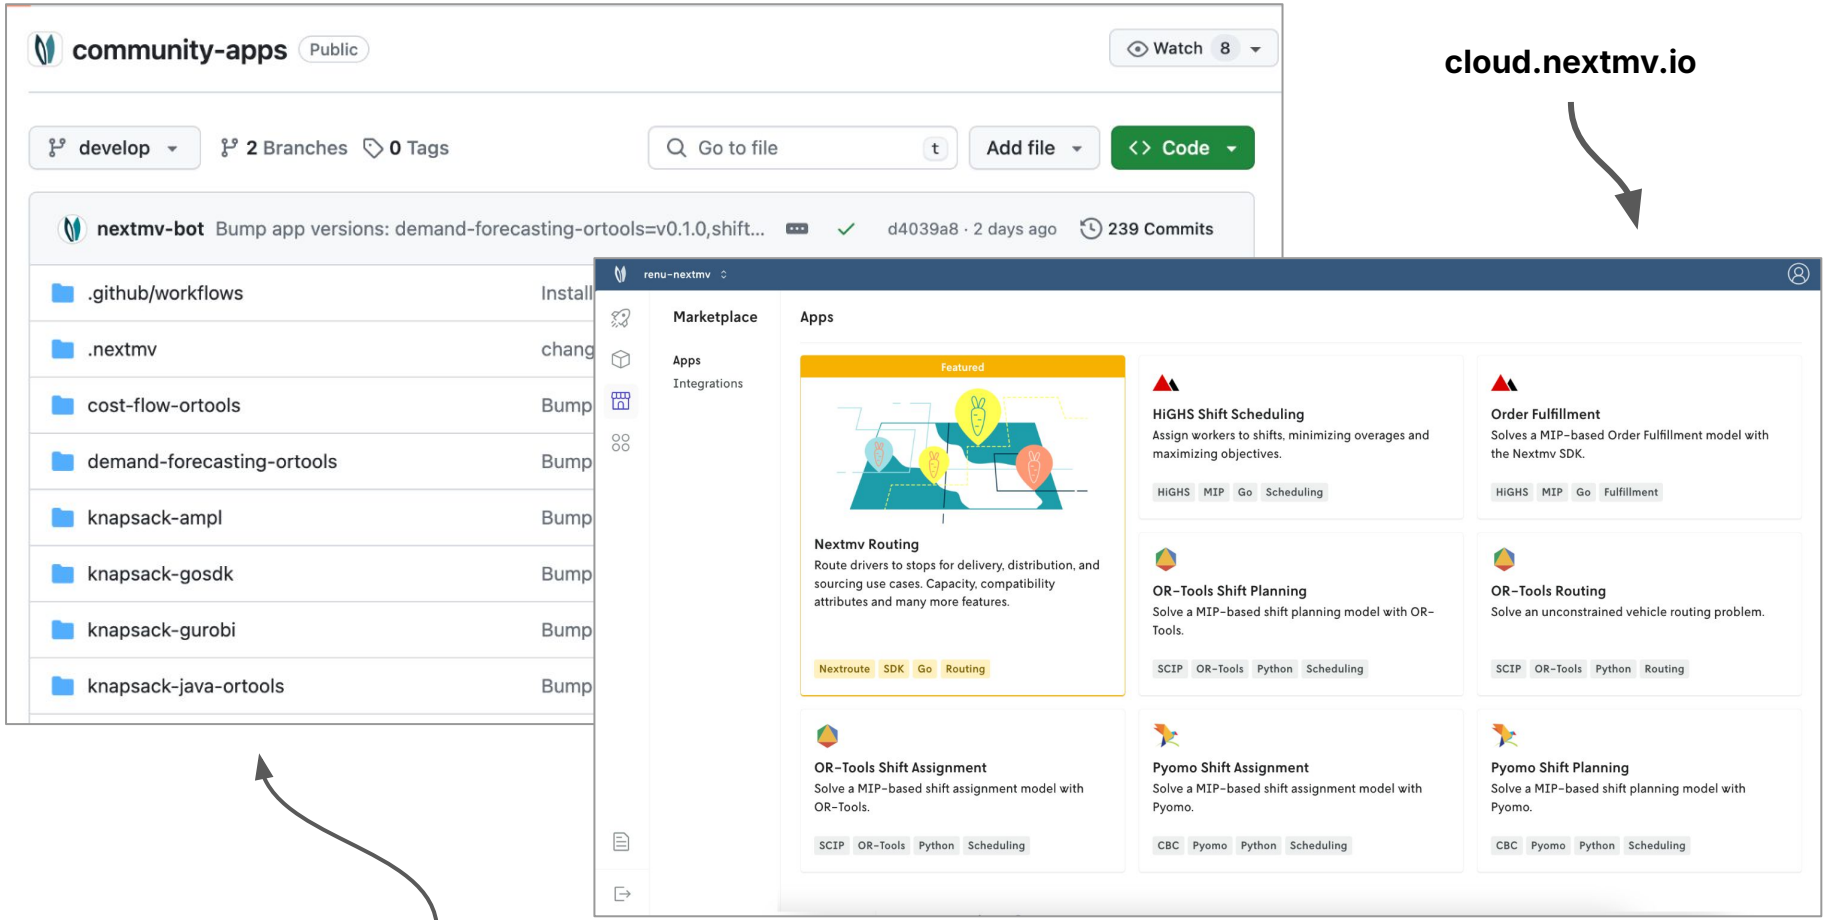

cloud.nextmv.io

https://github.com/nextmv-io/community-apps

This is an exciting time!

🥧 Our challenge is one of adoption. Solvers can lean into strengths instead of competing as commodities. If we grow the pie of optimization users, we all get more pie.

🏎️ Just like DS and ML, we need lots of “on ramps” to adoption. This means high quality commercial and open source solvers and platforms.

📈 The release of HiGHS marked an inflection point in open source solver development. It is accelerating!

... and so is commercial solver development!

https://www.linkedin.com/in/ryanjoneil/

linkedin.com/in/ryanjoneil/

@ryanjoneil

www.nextmv.io

Time for Q&A


nextmv