Building Pipelines and Deploying Models with *targets*, *tidymodels,* and *vetiver* in R

James Bristow

Junqi Zhu

R Exchange 2026

7th May 2026

Motivation

- Messy statistical and machine learning workflows
  - Ad-hoc scripts and notebooks
  - Not reproducible
  - Model deployment is an afterthought
- Solution: Tidy things up
  - Formalise workflows within pipelines
  - Standardise and modularise code into reusable components
  - Treat models as versioned and deployable S3 artifacts with documentation

DevOps and MLOps

- DevOps
  - Practices unifying software development and IT operations
  - Incorporating automation, infrastructure-as-code, and monitoring
  - Releasing stable code into production reliably and efficiently
- MLOps
  - Incorporates DevOps principles for machine learning (ML) systems
  - Ideas: Data & model versioning, reproducible pipelines, model deployment and cards
  - Reproducible, scalable, and trustworthy ML systems

MLOps Pipeline in R

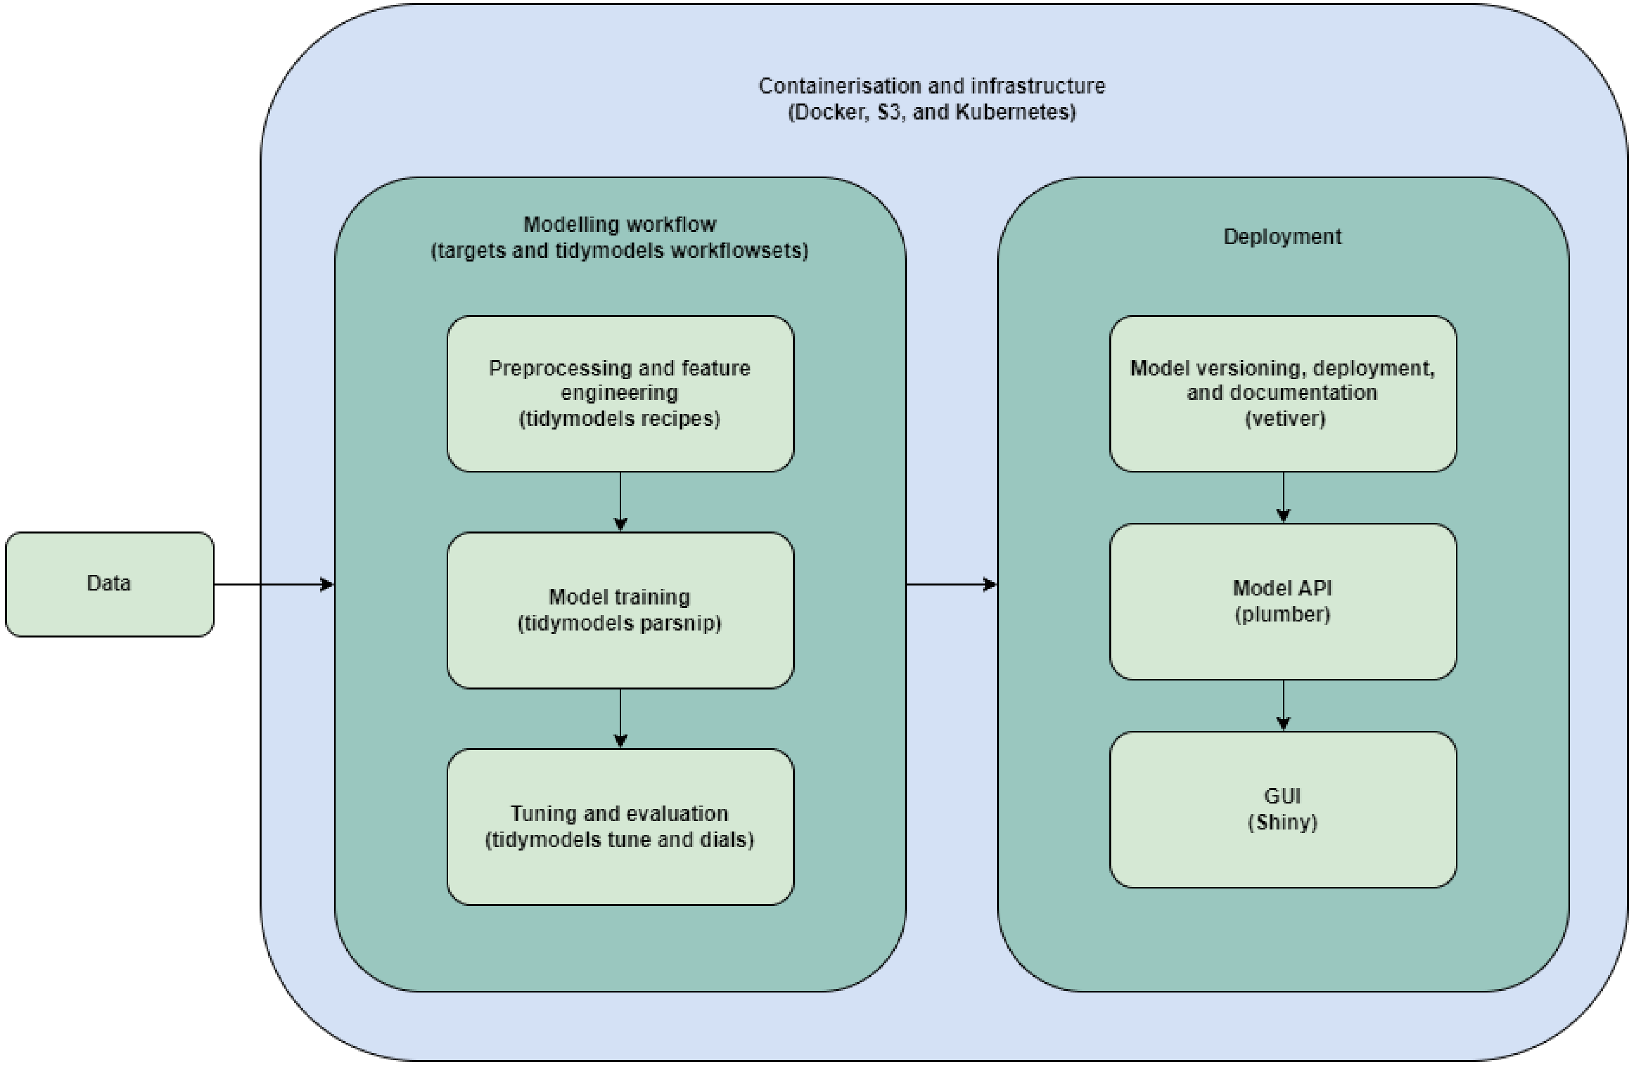

Case study: Grapevine yield prediction

- Marlborough wine
- Two vs four cane pruning
- Four different vineyards
- Predicting yield from climatic conditions
- Predictive ML model: Conformalised stacking ensemble

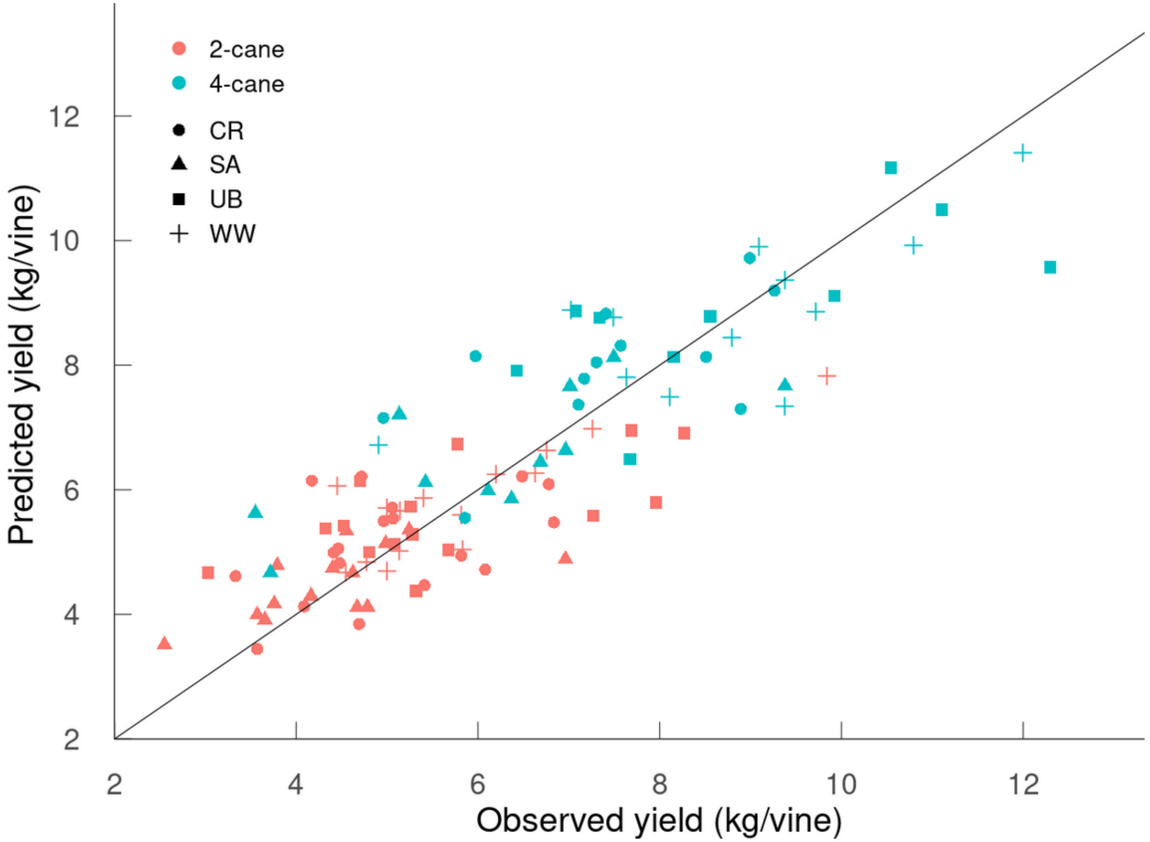

Reproducible targets pipelines

- Defines pipelines as a dependency graph (DAG)
- Automatically runs only what has changed

- Ensures end-to-end reproducibility, reducing manual reruns and "script chaos"
- Different backends via crew: Local; cloud (AWS, Google); and HPC (SLURM, SGE) execution environments

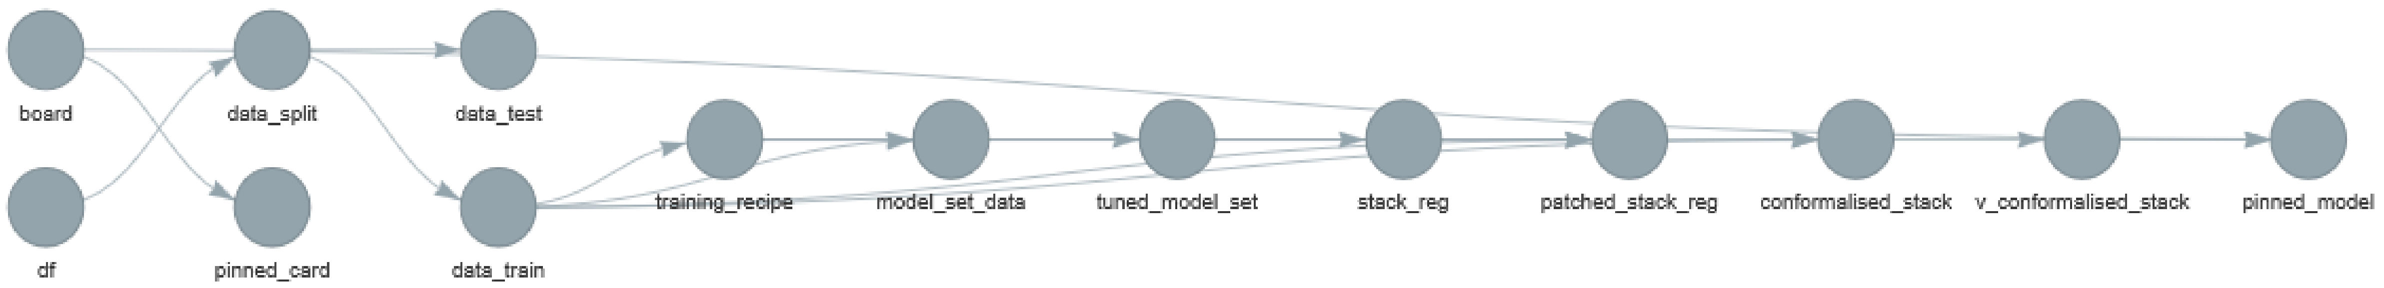

tidymodels

- Combines preprocessing, modelling, and tuning in one workflow
- Consistent interface across model types
- Fair comparison and systematic tuning
- Supports reproducible modelling pipelines
- tidyverse successor to caret
  - caret: Monolithic package – one interface
  - tidymodels: Modular – swap and extend components

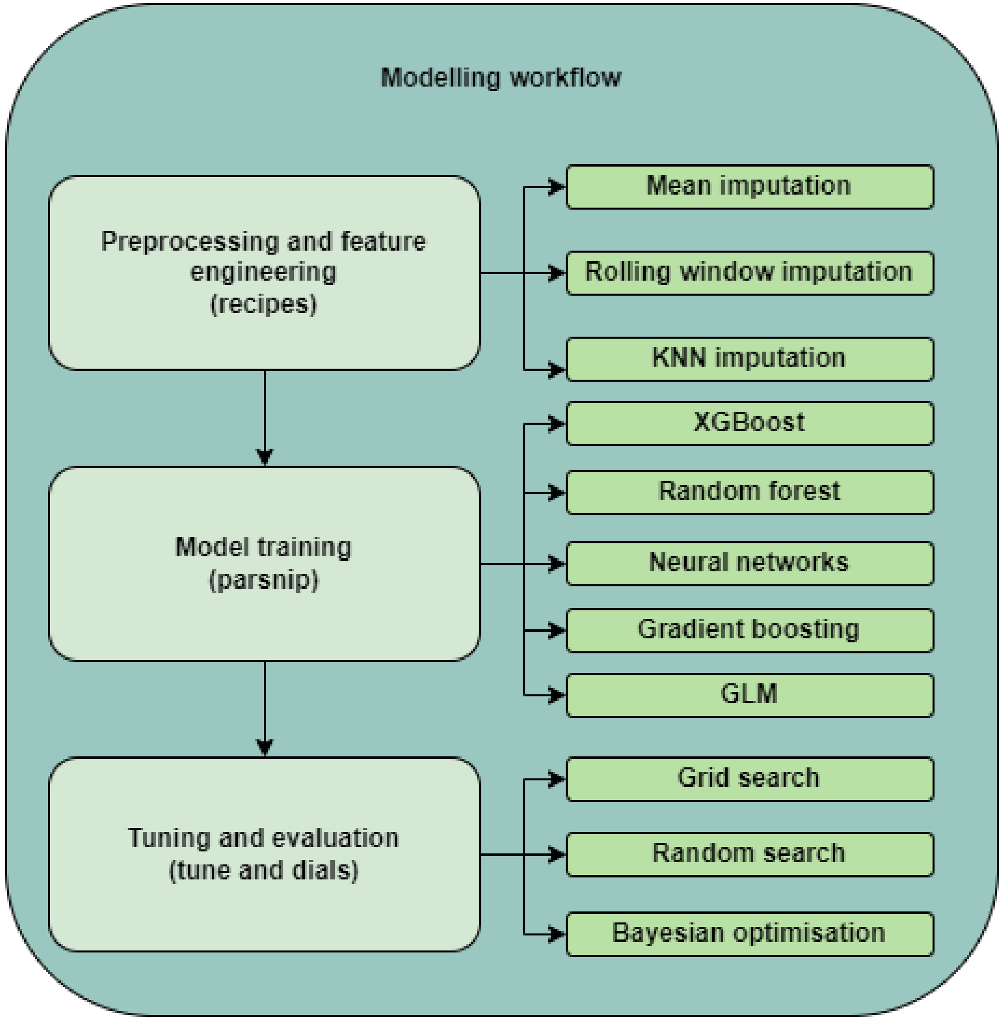

workflowsets and stacking ensembles

- **workflowsets**: Systematic generation of all preprocessing, models, and tuning combinations
- Efficient metric-based comparison and selection of workflows

	XGBoost	Random forest	Neural network	Gradient boosting	GLM
Mean imputation	XB1	RF1	NN1	GB1	GLM1
RW imputation	XB2	RF2	NN2	GB2	GLM2
KNN imputation	XB3	RF3	NN3	GB3	GLM3

X

Grid search	**Random search**	**Bayesian optimisation**

= 45 workflows

- Stack workflowset to improve predictive performance (1 linear metamodel)
- **stacks**: tidymodels package for creating stacking ensembles

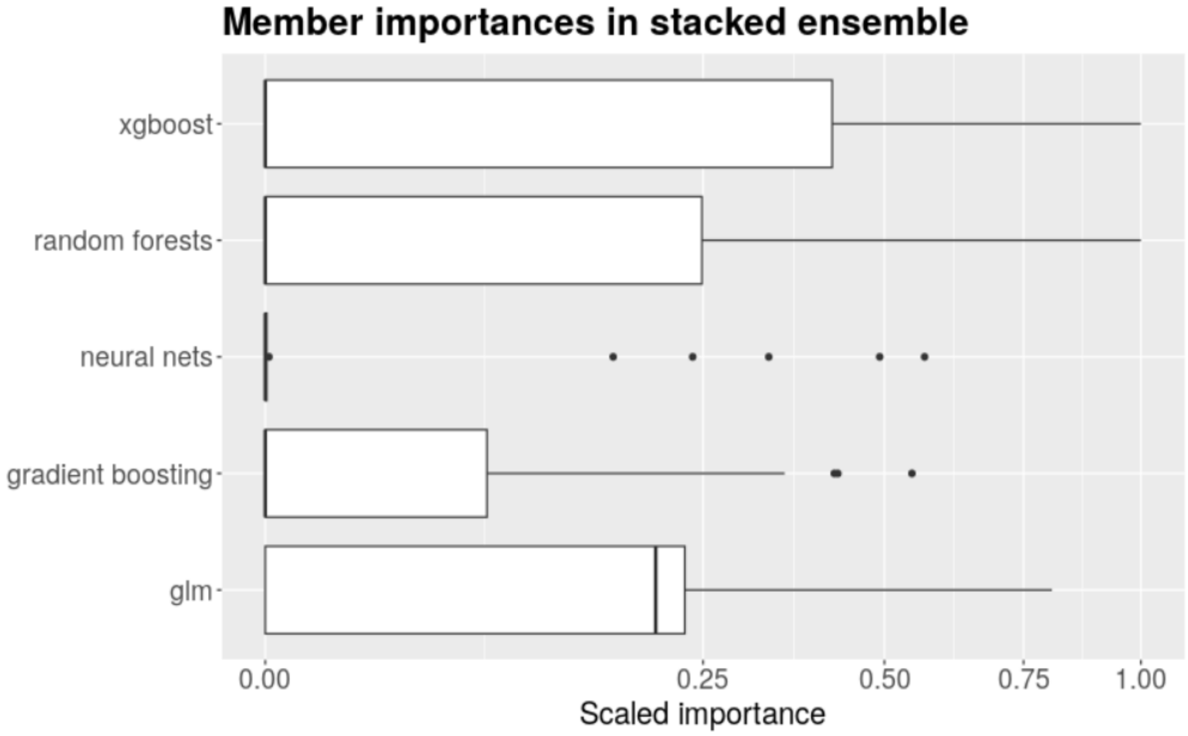

Conformal inference

- Model agnostic: Produce prediction intervals for any ML model
- Quantifies predictive uncertainty: Risk-sensitive applications
- Theoretically guaranteed coverage: e.g. 95% intervals contains truth ~95% of the time
- No distributional assumptions. Exchangeability assumption
- **probably:** tidymodels package
  for conformalisation and recalibration

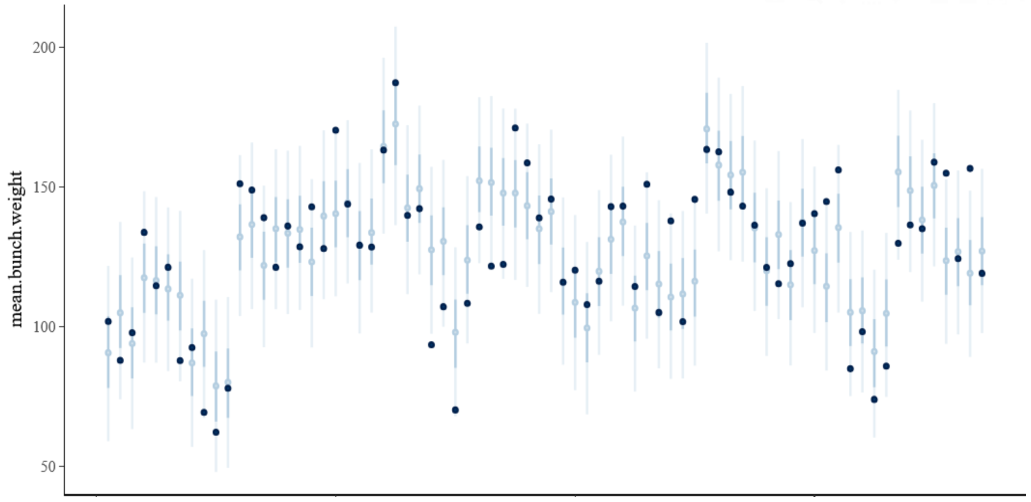

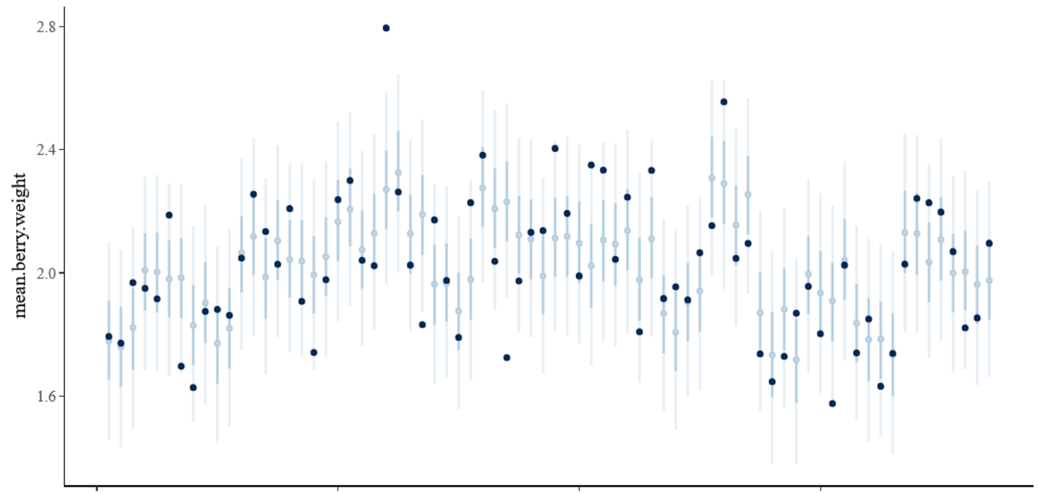

Model interpretability

- Makes black-box model decisions interpretable to humans
- Detect bias, leakage, or unexpected behaviour
- Debugging and model improvement
- Increases trust with stakeholders and decision-makers
- **DALEX**: Model interpretability toolbox; offers tidymodels integration

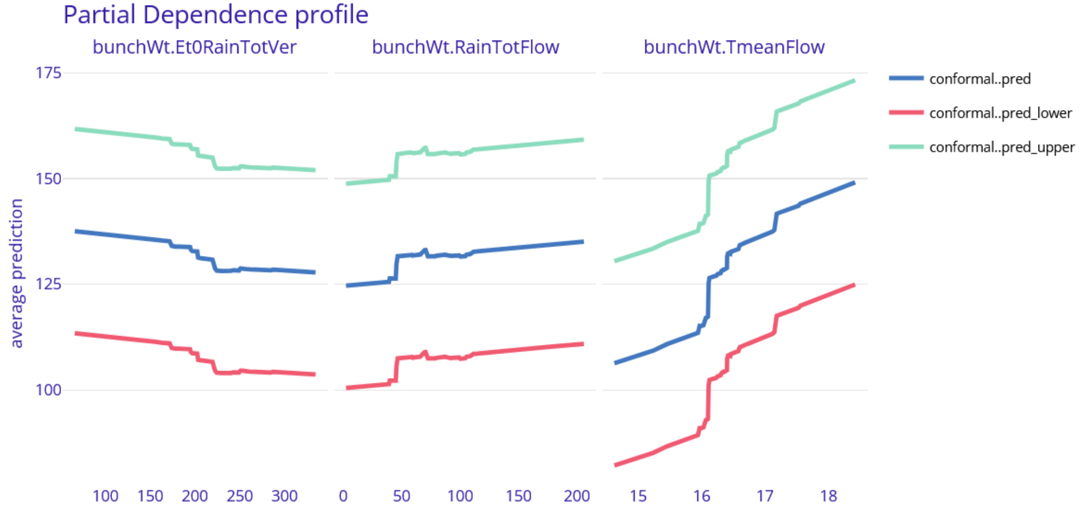

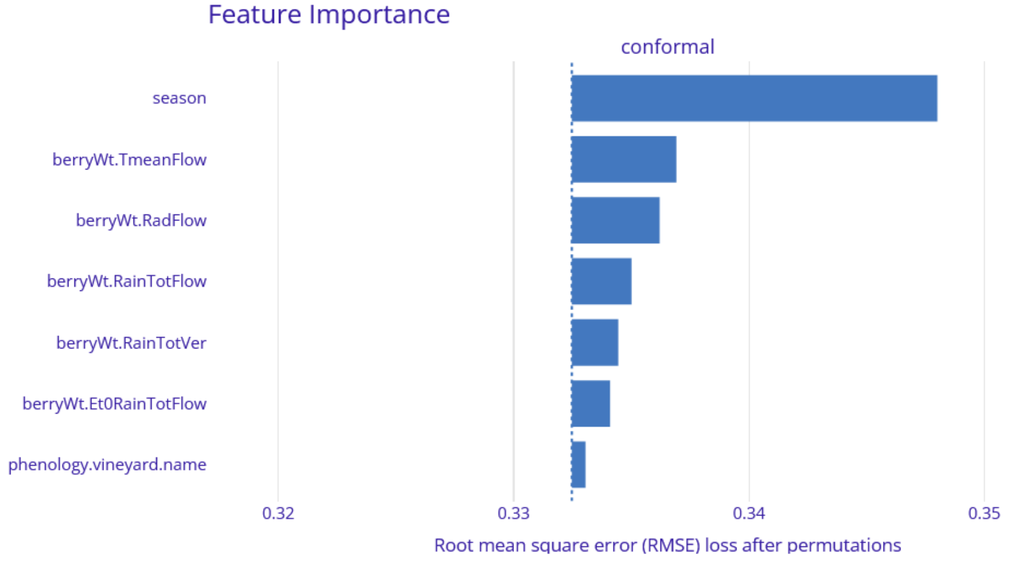

From ML model to API

- **vetiver:** Packages, version-controls, and deploys models
    - Consistent model-to-API process
    - Bridges gap between analysis and deployment
- **plumber:** Expose models as lightweight web apps (REST APIs)
    - Integration with apps, dashboards, and services

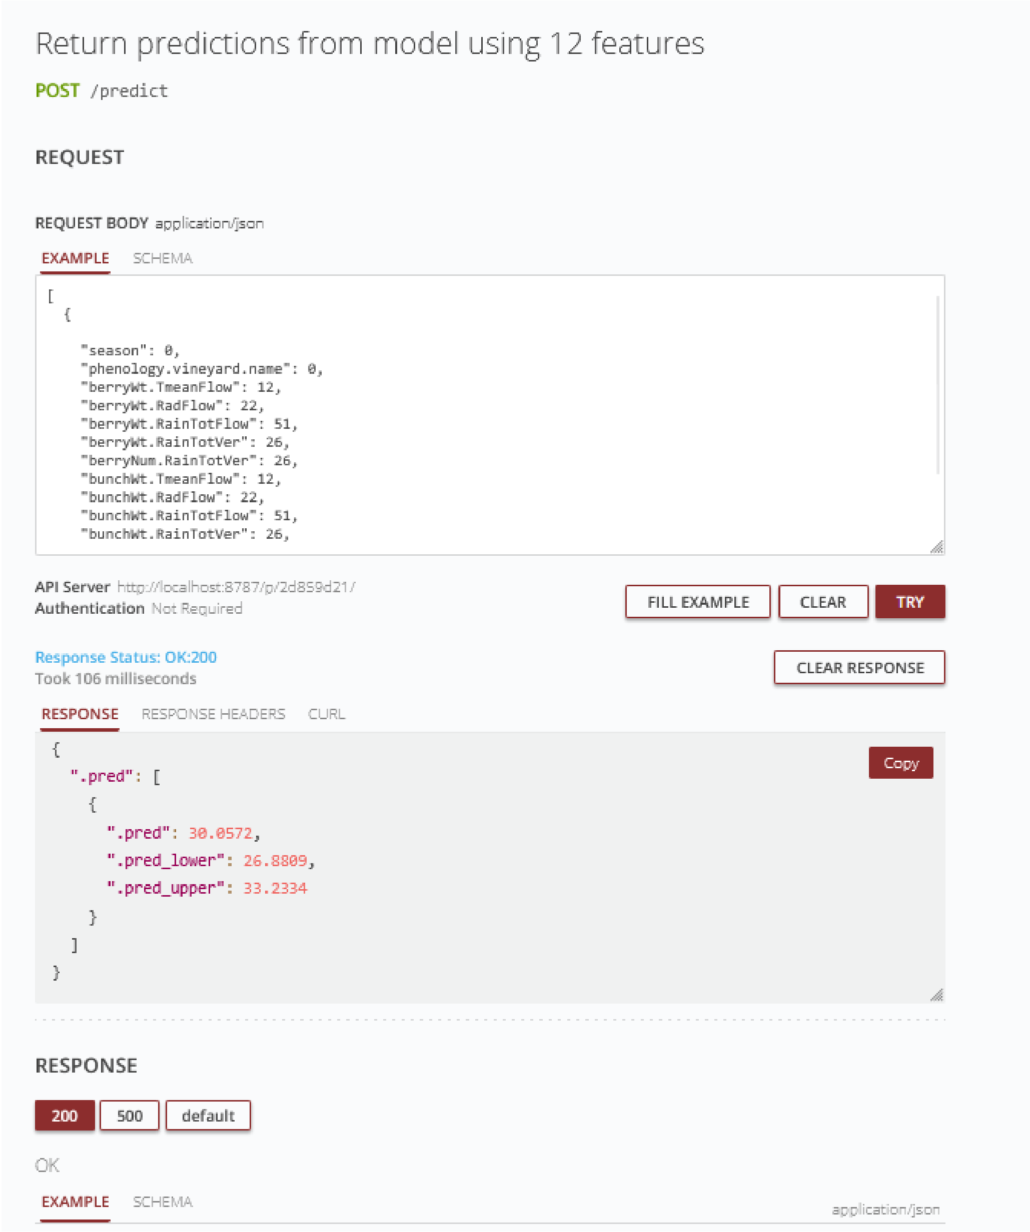

Dashboarding and model cards

- Shiny dashboards turn model outputs into interactive tools for decision-making
- vetiver model cards document model purpose, performance, and limitations

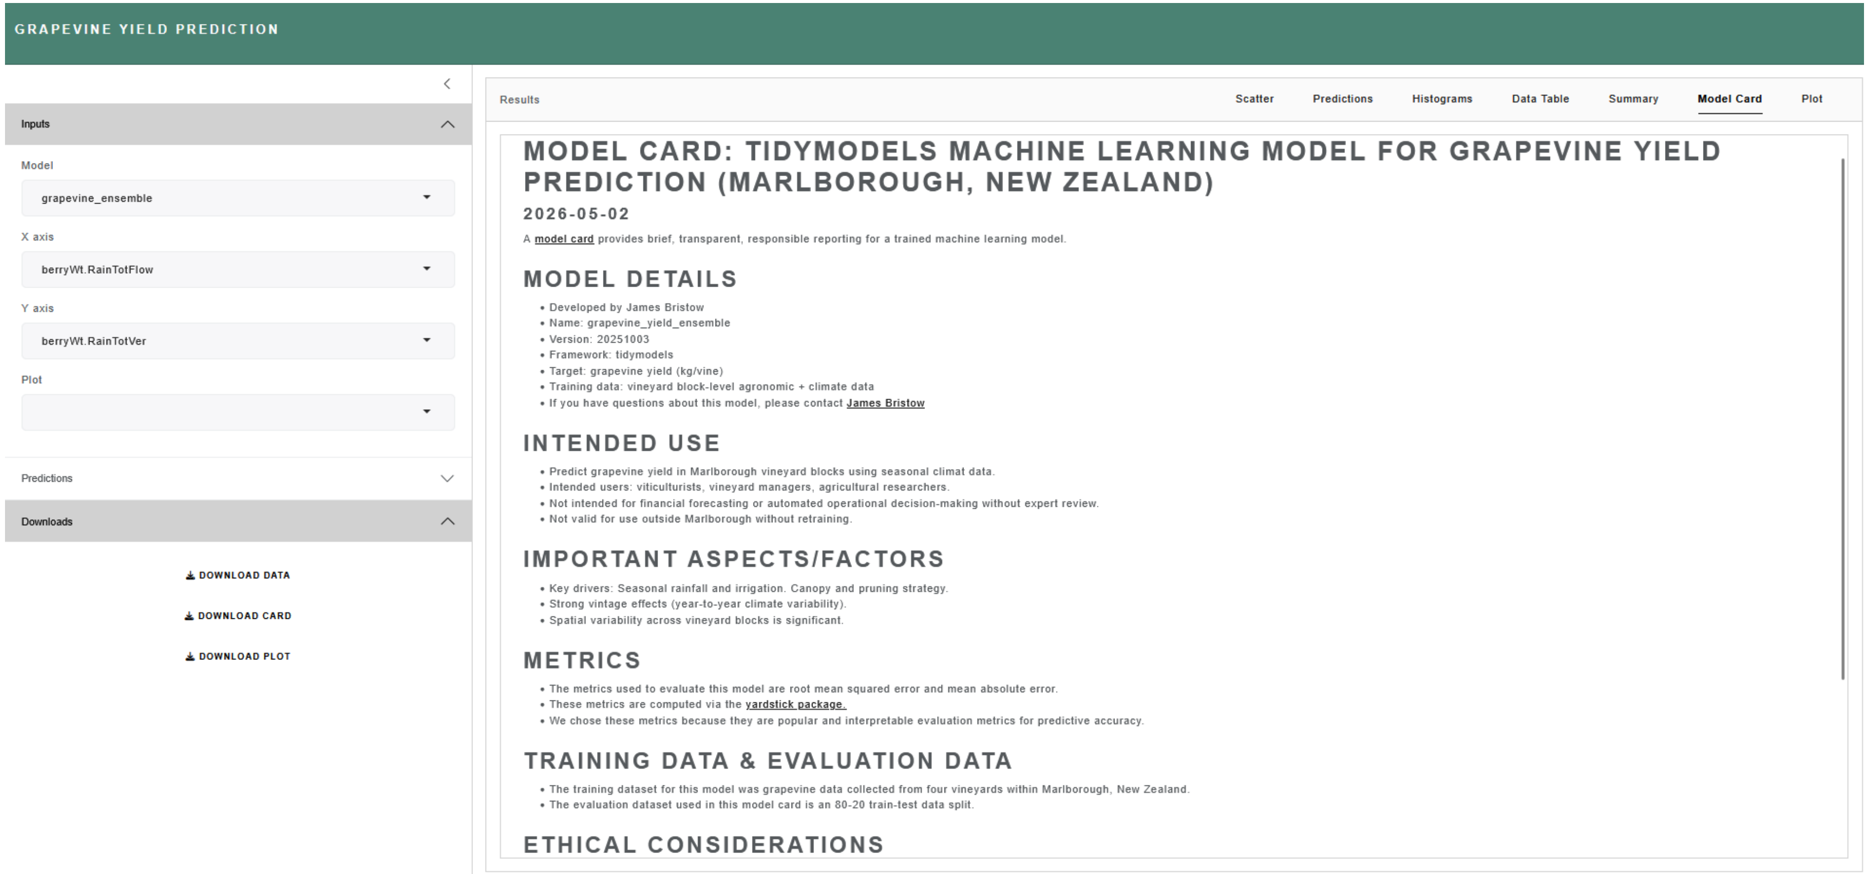

Containerisation and infrastructure

- S3: Storage
  - *Store it anywhere*
  - Store data, models, and artifacts centrally for reuse via pins

- Docker: Reproducibility
  - *Run it anywhere*
  - Package models and R environments for consistent execution

- Kubernetes: Scalability
  - *Run it everywhere*
  - Orchestrate deployments so models can scale and serve anywhere

Amazon S3

docker

kubernetes

Conclusion

- From ad-hoc scripts and notebooks to reproducible pipelines
- Adoption of DevOps principles within ML workflows
- Treat models as storable artifacts
- For more: https://github.com/Plant-Food-Research-Open/r-pipeline-development-workshop/tree/main

Good models matter.
But good pipelines decide whether they ever get used.

Acknowledgements

Funded by the New Zealand Ministry of Business, Innovation and Employment Endeavour Fund under the research project "Predicting the unseen: A new method for accurate yield estimation in viticulture/horticulture" for The University of Canterbury (UC) and The New Zealand Institute for Bioeconomy Science Limited (BSI).

BSI Staff who are working in the development of these platforms and those who have supported and worked in the programmes leading to this work. In particular, Junqi Zhu for his supervision during this work.

James.Bristow@plantandfood.co.nz

bioeconomyscience.co.nz

Thank you

Presentation disclaimer

Presentation for
R Exchange 2026, Wellington, New Zealand, 07 May 2026

Publication data:
James Bristow, Junqi Zhu, May, 2026, Building Pipelines and Deploying Models with *targets*, *tidymodels,* and *vetiver* in R. A PowerPoint presentation.
Job Code: P/471029/01 , SPTS No. 28537.

Presentation prepared by:
James Bristow
Data Scientist, Data Applications
May 2026

Presentation approved by:
Mark Wohlers
Science Group Leader, Data Science
May 2026

For more information contact:
James Bristow
Email: James.Bristow@plantandfood.co.nz

This report has been prepared by New Zealand Institute for Bioeconomy Science Limited (Bioeconomy Science Institute).
120 Mt Albert Road, Mt Albert, Auckland 1025, New Zealand, Tel: +64 9 925 7000, Fax: +64 9 925 7001.
www.plantandfood.co.nz