


# Advanced Machine Learning with Tidymodels

**Level:** Advanced  
**Duration:** 6 hours

A course that builds on the material covered in our [Machine Learning with Tidymodels](#) course. We take a look at how we can fit linear discriminant analysis (LDA) models using `{discrim}`, assessing model reliability using V-fold cross validation, pre-processing, tree-based models & more. If you wish to explore the abundance of model fitting techniques `{tidymodels}` has to offer, then this course is certainly for you!



## Course Outline

- **Qualitative modelling:** An introduction to Linear Discriminant Analysis (LDA).
- **Fitting LDA models & model prediction:** Using `{discrim}` to add additional bindings to `{parsnip}`, to fit LDA models
- **Advanced model assessment:** V-fold cross validation using `{workflows}` and `{tune}`
- **Penalised regression techniques:** Ridge regression, lasso and elastic net
- **Tree-Based Models:** Bagging, random forests and boosting

# Learning Outcomes

## Session 1:

*By the end of session 1, participants will be able to...*

- understand and fit linear discriminant analysis models in the Tidymodels style using {discrim} and {parsnip}.
- use {yardstick} to assess models and gain an understanding of model metrics.
- understand V-fold cross validation and bootstrapping and how to use it on your training data.
- tune model hyperparameters in an intelligent fashion using the {tune} and {dials} packages.

## Session 2:

*By the end of session 2, participants will be able to...*

- understand, create and tune penalised regression models and explore multiple metrics such as Ridge, Lasso and Elastic Net regression.
- understand variable importance and know how to use {vip} to assess variables.
- understand the concept of tree-based models and their uses in regression and classification.
- refine the performance of tree-based models by exploring ensembles: bagging, the {baguette} package, and random forest strategies.
- fit, tune and assess the predictive performance of tree-based models.

*This course does not include:*

- Other clustering alternatives to the K-nearest neighbours method.
- Dimensionality reduction methods except for linear discriminant analysis.
- Bayesian networks, neural networks and genetic algorithms.
- Computer Vision and Image recognition analysis.
- Boosting ensembles for tree-based models.

## Prior Knowledge

It will be assumed that participants are familiar with R. For example inputting data, basic visualisation, basic data structures and use of functions. In addition, attendees should be familiar with the concepts covered in our [Machine Learning with Tidymodels](#) course.

## Attendee Feedback

- "Very well taught. I will certainly be recommending this course to my colleagues!"

## Contact

[hello@jumpingrivers.com](mailto:hello@jumpingrivers.com)