



Deep Neural Networks for Regression Problems: From Data to Deployable Julia Models

Dr Jan Blomerus & Franco Deacon (*both at the University of the Free State*)

Attend 4 online lectures:

Wednesday 1, 15, 22 & 29 July 2026 | 14h00-15h00

ABSTRACT

This four-week series introduces the practical development of deep feed-forward neural networks for regression problems, with a focus on structured data and applications in Julia. Participants will be guided through the full modelling workflow: from data preparation and exploratory analysis to neural network design, training, validation, tuning, model comparison, and final model selection.

The sessions will include demonstrations using app.samplingforml.edu.eu.org, an interactive tool for exploring sampling in machine learning, and training DNNs for regression effectively. Where appropriate, the series will also draw on PCA and heatmap-based exploration to show how data structure, sampling decisions, and variable relationships can influence downstream model performance.

Emphasis will be placed on the practical modelling choices that matter, including architecture design, activation functions, loss functions, optimisation, and evaluation of predictive accuracy. The series is intended for participants interested in correctly applying neural networks to continuous-outcome prediction problems in a transparent and practically useful way.

LECTURES

LECTURE 1
1 July

Data, Sampling, and Problem Setup

Participants are introduced to regression problems involving structured data and continuous response variables. The session covers preprocessing, data splitting and exploratory analysis, with demonstrations using app.samplingforml.edu.eu.org for sampling, PCA, heatmap-based exploration, standardisation, and suggested baseline Python and Julia models.

LECTURE 2
15 July

Julia and Building a First DNN Model

This session introduces the Julia workflow for developing deep feed-forward neural networks for regression. Participants will move from prepared data to a first working and practically demonstrated DNN model, covering basic model structure, activation functions, loss functions, optimisation, and training.

LECTURE 3
22 July

Improving and Tuning the DNN

Participants will refine initial DNN models by exploring architecture choices, learning rates, batch sizes, and epochs, and by tuning hyperparameters. The focus is on improving predictive performance while understanding the effect of modelling decisions.

LECTURE 4
29 July

Evaluation, Comparison, and Final Model Choice

The final session considers how regression models are evaluated and compared. Participants will examine predictive performance measures, compare DNNs with traditional regression approaches, and discuss what matters and should be included in a final model and summary.



Jan Blomerus is a Senior Lecturer whose research focuses on actuarial science, statistical modelling and the application of deep learning to insurance, using neural networks to analyse complex datasets.



Franco Deacon is a final-year BSc Actuarial Science student at the University of the Free State with research interests in agentic AI, deep neural networks, actuarial science and finance.

REGISTER: <https://bit.ly/4f4UFEk>

