Introduction to Scala and functional programming

Level: Advanced

Duration: undefined hours

We are very happy to announce that Prof Darren Wilkinson is running a series of four courses for data science and statistics with Scala.

Course 1 will begin with an introduction to the Scala language and basic concepts of functional programming, as well as essential Scala tools such as Sbt for managing builds and library dependencies. A brief introduction to the IntelliJ IDE will also be provided. The main emphasis will be on the latest version of Scala, Scala 3, but Scala 2 will also be discussed. The course will continue with an overview of the Scala collections library, including parallel collections, and we will see how parallel collections enable trivial parallelisation of many algorithms on multi-core hardware.

Course Outline

Course 1 will begin with an introduction to the Scala language and basic concepts of functional programming (FP), as well as essential Scala tools such as Sbt for managing builds and library dependencies. A brief introduction to the IntelliJ IDE will also be provided. The main emphasis will be on the latest version of Scala, Scala 3, but Scala 2 will also be discussed. The course will continue with an overview of the Scala collections library, including parallel collections, and we will see how parallel collections enable trivial parallelisation of many algorithms on multi-core hardware.

This suite of 4 half-day courses is aimed at statisticians and data scientists already familiar with a dynamic programming language (such as R, Python or Octave). Scala is a free modern, powerful, strongly-typed, functional programming language. It is fast and efficient, runs on the Java virtual machine (JVM), and is designed to easily exploit modern multi-core and distributed computing architectures. Scala is a favourite language for data engineering teams and others wanting to work with data at scale in an efficient, safe and timely fashion. For similar reasons, it is also very well suited to the development of robust data science, machine learning and statistical applications.

The courses can be taken independently, but do have pre-requisites which are detailed within the Prior Knowledge summaries. They will be delivered through a combination of lectures, live demos and hands-on practical sessions. The courses will be delivered by Prof Darren Wilkinson, a well-known expert in computational Bayesian statistics and a leading proponent of the use of strongly-typed FP languages (such as Scala) for scalable statistical computing. Participants will be expected to use their own laptops and to have a recent version of Java pre-installed. Other set-up instructions will be provided in advance to registered participants.

Prior Knowledge

Course 1 assumes basic programming experience. It is assumed that participants will be familiar with writing their own functions in a language such as R or Python, including essential control structures such as 'for-loops' and 'if-statements'. The course is not suitable for people completely new to programming. However, no prior knowledge of Scala or functional programming is assumed. All participants will be expected to bring their own (multi-core) laptop and to have a recent version of Java pre-installed. Other set-up instructions will be provided in advance to registered participants.

Contact

hello@jumpingrivers.com