

Budapest University of Technology and Economics
Department of Measurement and Information Systems
Fault Tolerant Systems Research Group

Critical Architectures Laboratory (Spring 2016/17)

Monitoring Data Analysis

Imre Kocsis, based on the previous iteration of the lab by Ágnes Salánki
v0.7, April 2, 2017.

1 Introduction

Data-driven approaches have always been present in the design, evaluation and operation of critical systems; however, the relatively recent emergence of "data science" is rapidly making them much more popular, easy to perform and accessible for the average practitioner. The main goal of this lab is to give a rather bird's-eye overview of applying modern, general purpose data analysis approaches and tools to the data analysis of critical infrastructures.

More specifically, our intention is to give an example-based introduction to the following topics:

1. The common contemporary idioms of data analysis in general purpose scripting and programming languages; specifically, the "data frame " abstraction and the most common operations over data frames.
2. Notebook-based data analysis. (Interactive) notebooks constitute a relatively new, but very important addition to the list of generic data analysis styles (classically mostly REPL shells/IDEs, explicit workflows, fully interactive data exploration tools and BI-style dashboards).
3. Plotting using rich library support. In the not so distant days of the past, general purpose data visualization in the practice meant picking a potentially closed, dedicated tool from a rather small set - most probably Excel, Matlab or gnuplot. One could argue that visually, the results were a bit lacking. In contrast, we have now a very wide array of options; conceptually one of the main advances is that visualization is getting very well "integrated into" analytic scripting and programming.
4. Integrated (and relatively painless) access to Big Data capabilities. "Big Data" tools have originally emerged to enable relatively cheap (at least compared to data warehouses) distributed processing of massive amounts of "data at rest" - i.e. data that will not be updated or deleted from. However, this broke the convenient status quo of "everything that stores data speaks SQL"; for a while, analysis of Big Data meant that the "data scientist" had to essentially perform at least some algorithm development as well as code development for the Big Data platform at hand - to achieve functionality that has been rather trivial in classis SQL-based as well as statistical scripting settings. A welcome development is that this is changing, too; in addition to Big Data tools gaining true SQL support, in a number of important cases access to Big Data processing is now very well integrated into "workstation-based" statistical scripting and programming.

2 Python or R?

The development of data science is now at a point where all our goals could be demonstrated on wide range of platforms and using a wide range of contemporary scripting and programming languages (e.g. manifestations of Data Frames are even present in C# and Java now). As our data analysis examples are geared towards data exploration, we focus on scripting languages.

Instead of specifying a language, this year the lab has two "tracks" in the sense that you are allowed to choose between Python and R, based on your existing experience. That being said, if you are not familiar with R, you are strongly suggested to do the lab preparation and the lab itself in Python.

3 Preparation for the lab

As a preparation for the lab, you are expected to create an account for the free service <https://datascientistworkbench.com/> (BDU Labs); we plan to use this platform. We want to emphasize that - with some minor exceptions - the tools that are bundled together in this Software as a Service solution are all free and open source; if you are so inclined, you are welcome to roll

and bring along your own environment (e.g. using Docker Hub images this should not be particularly hard).

Additionally, you should familiarize yourself with the following material. We do not expect full proficiency; however, to finish the lab assignment in a timely manner you *will* have to know where to look and what to search for.

3.1 Jupyter/IPython

On BDU Labs, perform the "The Basics" tutorials. The examples are for Python, but with some existing experience in R can be easily translated for the R kernel. Note that while R is an option, RStudio, although it is present in the SaaS bundle, is not; you will be expected to work in Jupyter (with the R kernel).

3.2 Pandas

From the official pandas documentation: <http://pandas.pydata.org/pandas-docs/stable/tutorials.html>, review the "10 minute tutorial" then proceed to lessons 1 through 6 from "Lessons for New pandas Users" (you can skip the Excel-specific parts). Reading the Cookbook may be advantageous, but is not essential. On the R side: as data frames are a language-level concept in R, if you know R, you should know at least one way to perform what the referred pandas material talks about. That being said, you can benefit from looking at the package dplyr and/or data.tables.

3.3 Pandas visualization (matplotlib with training wheels)

With matplotlib under the hood, Pandas has decent built-in visualization support - and this should be enough for the lab assignments. Note that while you are welcome to use matplotlib more directly (or use the Python visualization library of your choice), struggling with the peculiarities of a visualization library can become a time sink very fast in any language.

You should review the respective part of the pandas documentation: <http://pandas.pydata.org/pandas-docs/version/0.19.0/visualization.html> Note: you are *not* expected to read and know this in detail; however, you should be aware of your options (what can be done and what can't). For R: ggplot2 is always a safe bet; base graphics will be frowned upon.

Note: due to your repeated exposure in the courses of the research group, we will expect you to know the basics (philosophy, approach, typical plot types) of Exploratory Data Analysis (EDA). You can find some Hungarian reference material from our group here:

<https://inf.mit.bme.hu/sites/default/files/materials/category/kateg%C3%B3ria/oktat%C3%A1s/bsc-t%C3%A1rgyak/rendszermodellez%C3%A9s/16/08-vizualis-elemzes.pdf>

and here:

http://www.interkonyv.hu/konyvek/antal_peter_intelligens_adatelemzes

Note that both sources approach EDA with a heavy emphasis on its "visual analysis" aspect. (This is due to the fact that your curriculum does not include statistics - so discussing the role and place of EDA in the statistical workflow in detail would lack the necessary underpinnings.)

3.4 Notebook-integrated Big Data: Spark

The lab will include a "Big Data" assignment - infrastructure monitoring data is a prime example of (valuable) "at rest" data for which bulk storage and distributed processing on "clusters of grey boxes" makes the most sense.

We will use one of the currently most successful Big Data platforms, Apache Spark. However, neither the scope nor the time constraints of the lab make an in-depth treatment of the subject possible. What we aim at is demonstrating that a) when offered in a Platform as a Service style,

Big Data services can be integrated into an application in a rather "painless" way (i.e. no need to acquire computing and storage resources, setting up the platform and configuring it); b) Big Data services have good language integrations that enable "ingesting" them in an analytic script easily. Consequently, while you are advised to gain familiarity with Spark *on the level of basic concepts*: <http://spark.apache.org/talks/overview.pdf>, we do not expect any deep proficiency. Spark has good (data frame) APIs that we will program against; review the "Spark SQL, DataFrames and Datasets Guide": <https://spark.apache.org/docs/latest/sql-programming-guide.html#getting-started> up to the Data Sources part. The lab assignment document will help you with getting around the specifics.

4 Self-check questions

If you are well-prepared for the lab, you should be able to answer the following questions:

1. From the point of view of data analysis workflow, how does the "Notebook" abstraction of Jupyter/IPython compare to classic tools running a REPL shell? (MATLAB, standard R, standard interactive command line Python, ...). What are the main benefits of Notebooks?
2. What is a data frame? (Either a) in general, b) specifically in pandas or c) specifically in R).
3. What are the main ways to filter a data frame in pandas/in R for a set of rows/set of columns?
4. How would you perform the "group by" operation on a data frame in pandas/R?
5. Specify a way for appending a new column to a data frame in pandas/R!
6. How are missing values represented in pandas/R?
7. What are the fundamental visualizations for a single observed variable? (Recall that data can be categorical as well as numeric).
8. What are the fundamental visualizations for two observed variables? (Recall that data can be categorical as well as numeric).
9. What is an RDD in Spark?
10. What are the most important Spark RDD Transformations?
11. What are the most important Spark RDD Actions?
12. Describe the difference between exploratory and confirmatory data analysis!