Data Science

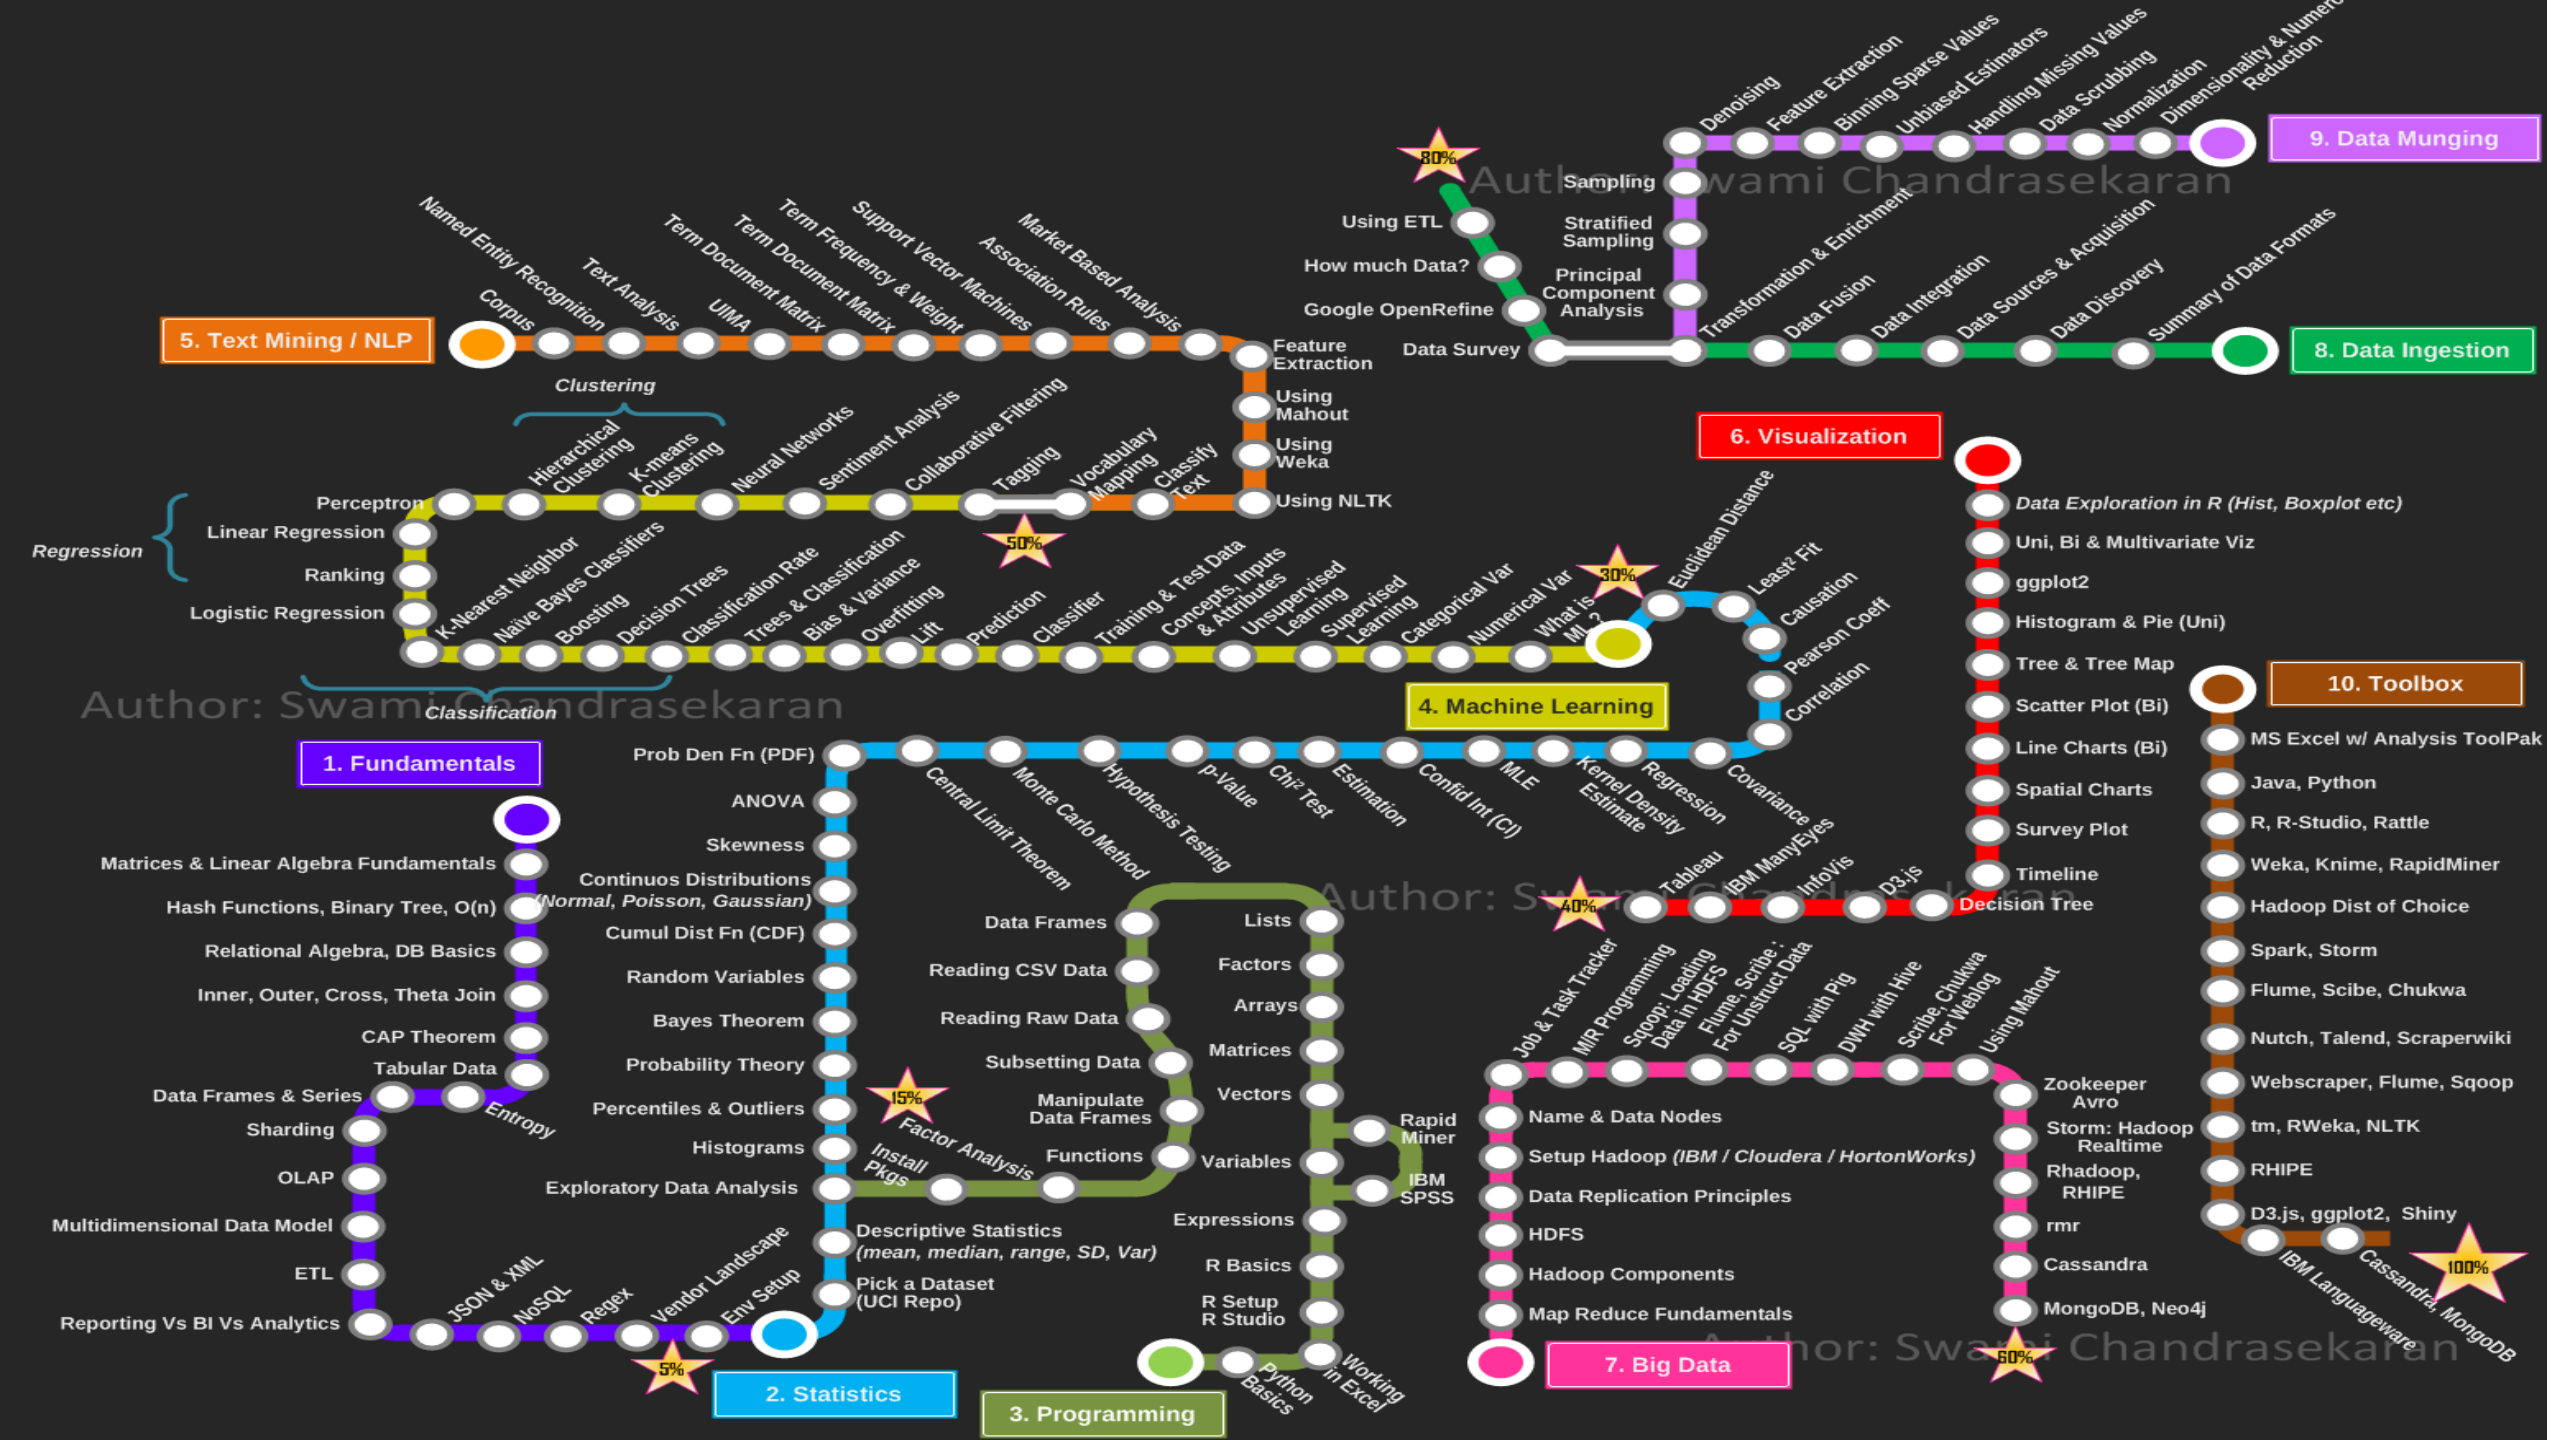
1. Fundamentals
Matrices & Linear Algebra Fundamentals
Hash Functions, Binary Tree, O(n)
Relational Algebra, DB Basics
Inner, Outer, Cross, Theta Join
CAP Theorem
Tabular Data
Data Frames & Series
Entropy
Sharding
OLAP
Multidimensional Data Model
ETL
Reporting Vs BI Vs Analytics
JSON & XML
NoSQL
Regex
Vendor Landscape
Env Setup
5%
2. Statistics
Prob Den Fn (PDF)
ANOVA
Skewness
Continuos Distributions (Normal, Poisson, Gaussian)
Cumul Dist Fn (CDF)
Random Variables
Bayes Theorem
Probability Theory
Percentiles & Outliers
Histograms
Exploratory Data Analysis
Descriptive Statistics (mean, median, range, SD, Var)
Pick a Dataset (UCI Repo)
Central Limit Theorem
Monte Carlo Method
Hypothesis Testing
p-Value
Chi² Test
Estimation
Confid Int (CI)
MLE
Kernel Density Estimate
Regression
Covariance
15%
Factor Analysis
Install Pkgs
3. Programming
Data Frames
Reading CSV Data
Reading Raw Data
Subsetting Data
Manipulate Data Frames
Functions
Lists
Factors
Arrays
Matrices
Vectors
Variables
Expressions
R Basics
R Setup
R Studio
Rapid Miner
IBM SPSS
Python Basics
Working in Excel
4. Machine Learning
Regression
Perceptron
Linear Regression
Ranking
Logistic Regression
Classification
K-Nearest Neighbor
Naive Bayes Classifiers
Boosting
Decision Trees
Classification Rate
Trees & Classification
Bias & Variance
Overfitting
Lift
Prediction
Classifier
Training & Test Data
Concepts, Inputs & Attributes
Unsupervised Learning
Supervised Learning
Categorical Var
Numerical Var
What is ML?
30%
Euclidean Distance
Least² Fit
Causation
Pearson Coeff
Correlation
Clustering
Hierarchical Clustering
K-means Clustering
Neural Networks
Sentiment Analysis
Collaborative Filtering
Tagging
Vocabulary Mapping
Classify Text
50%
Using NLTK
Using Weka
Using Mahout
Feature Extraction
Named Entity Recognition
Text Analysis
UIMA
Term Document Matrix
Term Document Matrix
Term Frequency & Weight
Support Vector Machines
Association Rules
Market Based Analysis
Corpus
5. Text Mining / NLP
6. Visualization
Data Exploration in R (Hist, Boxplot etc)
Uni, Bi & Multivariate Viz
ggplot2
Histogram & Pie (Uni)
Tree & Tree Map
Scatter Plot (Bi)
Line Charts (Bi)
Spatial Charts
Survey Plot
Timeline
Decision Tree
40%
Tableau
IBM ManyEyes
InfoVis
D3.js
7. Big Data
Job & Task Tracker
MR Programming
Sqoop: Loading Data in HDFS
Flume, Scribe : For Unstruct Data
SQL with Pig
DWH with Hive
Scribe, Chukwa For Weblog
Using Mahout
Name & Data Nodes
Setup Hadoop (IBM / Cloudera / HortonWorks)
Data Replication Principles
HDFS
Hadoop Components
Map Reduce Fundamentals
Zookeeper Avro
Storm: Hadoop Realtime
Rhadoop, RHIPE
rmr
Cassandra
MongoDB, Neo4j
60%
8. Data Ingestion
Summary of Data Formats
Data Discovery
Data Sources & Acquisition
Data Integration
Data Fusion
Transformation & Enrichment
Data Survey
Principal Component Analysis
Google OpenRefine
How much Data?
Using ETL
80%
Sampling
Stratified Sampling
9. Data Munging
Denoising
Feature Extraction
Binning Sparse Values
Unbiased Estimators
Handling Missing Values
Data Scrubbing
Normalization
Dimensionality & Numerosity Reduction
10. Toolbox
MS Excel w/ Analysis ToolPak
Java, Python
R, R-Studio, Rattle
Weka, Knime, RapidMiner
Hadoop Dist of Choice
Spark, Storm
Flume, Scibe, Chukwa
Nutch, Talend, Scraperwiki
Webscraper, Flume, Sqoop
tm, RWeka, NLTK
RHIPE
D3.js, ggplot2, Shiny
IBM Languageware
Cassandra, MongoDB
100%
Author: Swami Chandrasekaran

Data Science Tool Usage Survey (2014/O'Rielly)

- Still dominated by simple tools…

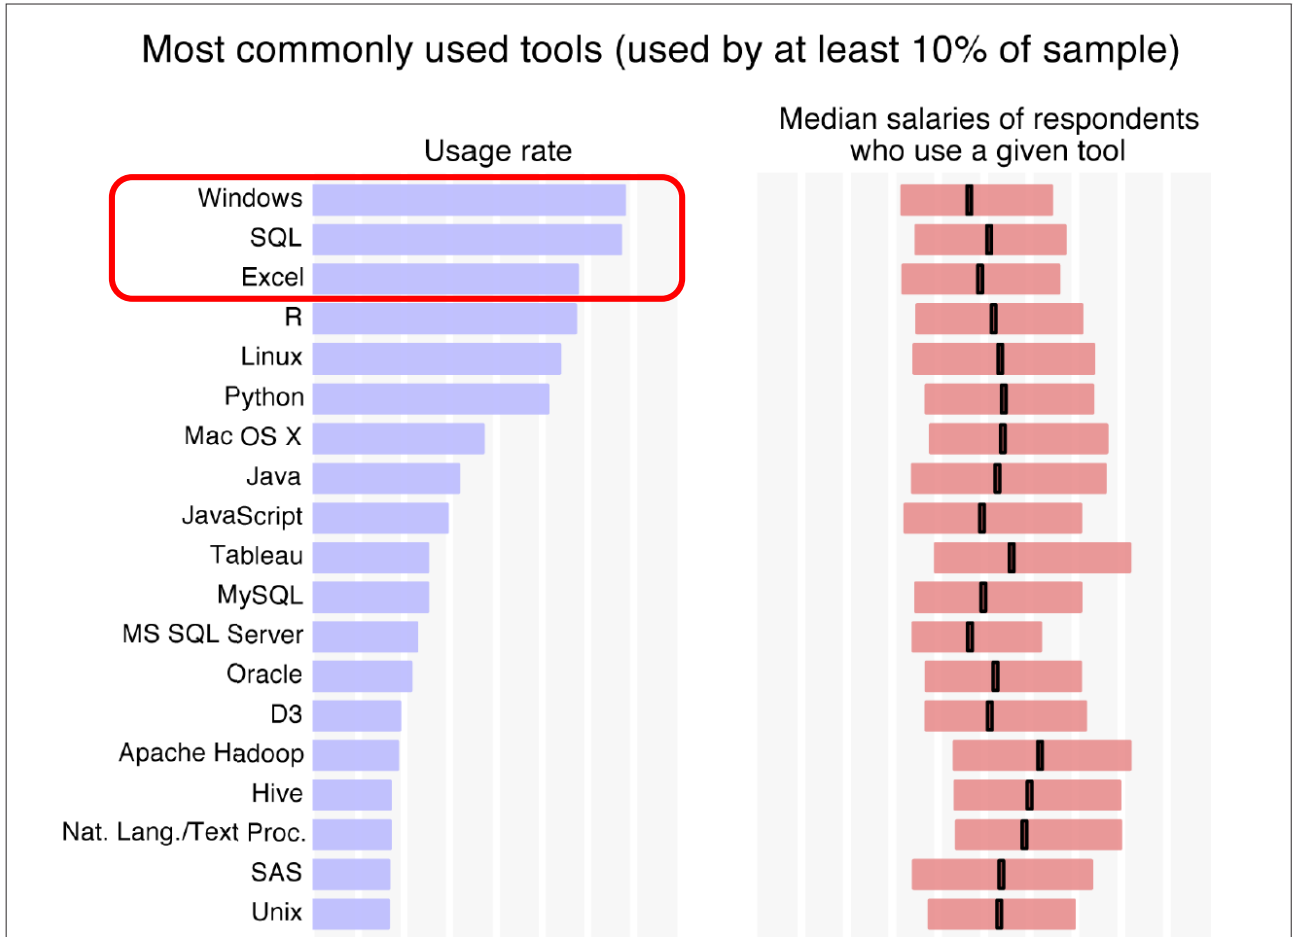

Choosing Tools for Data Science

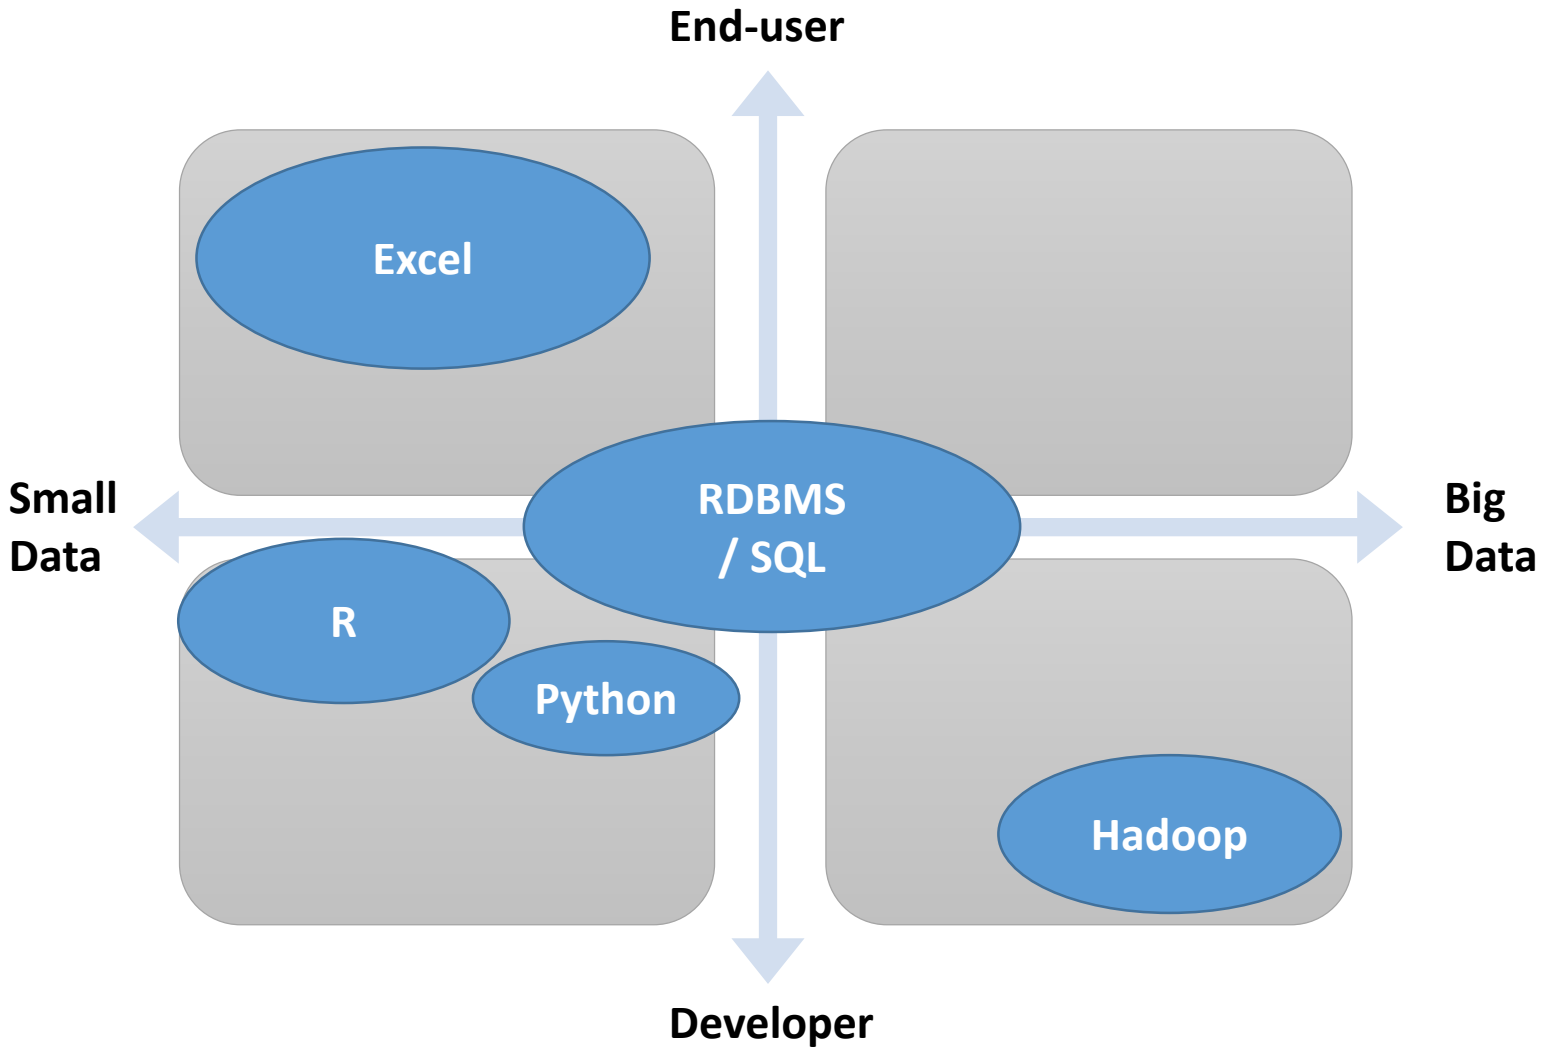

Chaining Tools for Data Science

- Use the right toolset in different stages

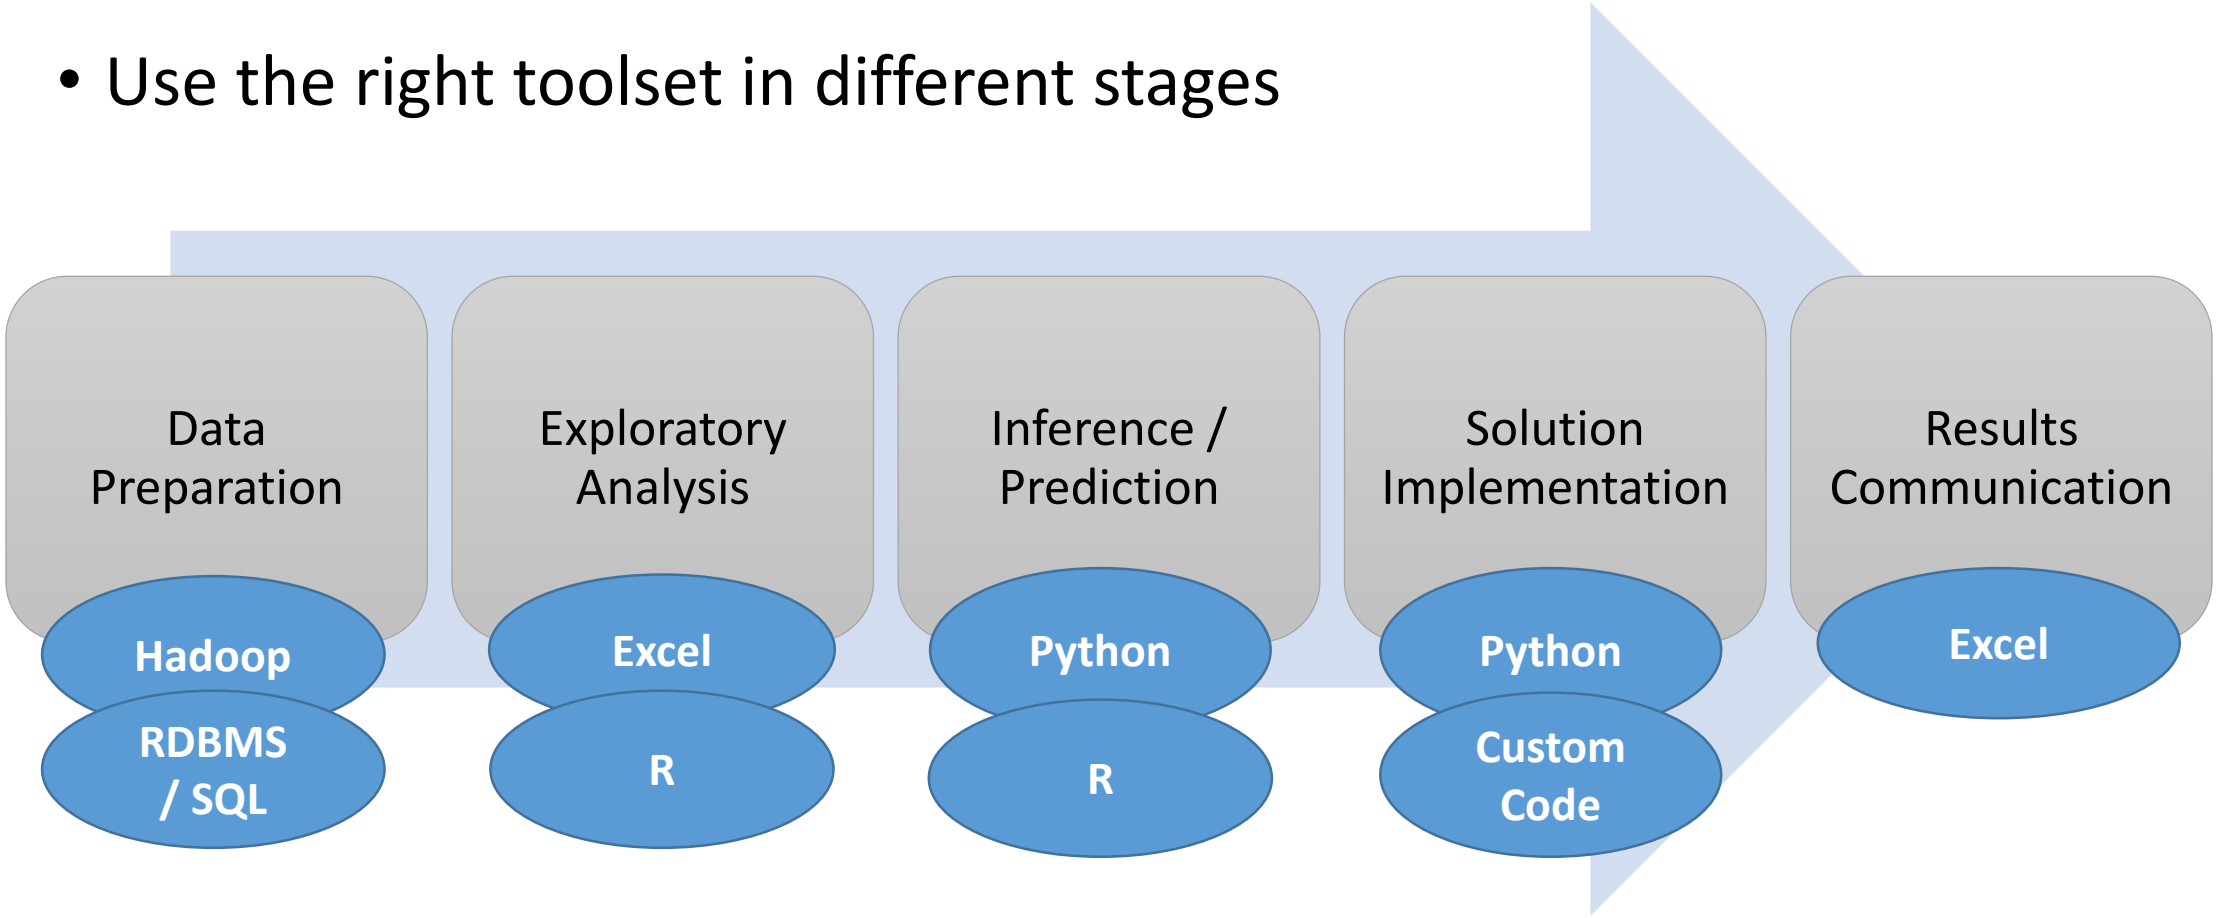

What is Data Science?

- An area that manages, manipulates, extracts, and interprets knowledge from tremendous amount of data
- Data science (DS) is a multidisciplinary field of study with goal to address the challenges in big data
- Data science principles apply to all data – big and small

https://hbr.org/2012/10/data-scientist-the-sexiest-job-of-the-21st-century/

What is Data Science?

- Theories and techniques from many fields and disciplines are used to investigate and analyze a large amount of data to help decision makers in many industries such as science, engineering, economics, politics, finance, and education
  - Computer Science
    - Pattern recognition, visualization, data warehousing, High performance computing, Databases, AI
  - Mathematics
    - Mathematical Modeling
  - Statistics
    - Statistical and Stochastic modeling, Probability.

Data Science Is Multidisciplinary

By Brendan Tierney, 2012

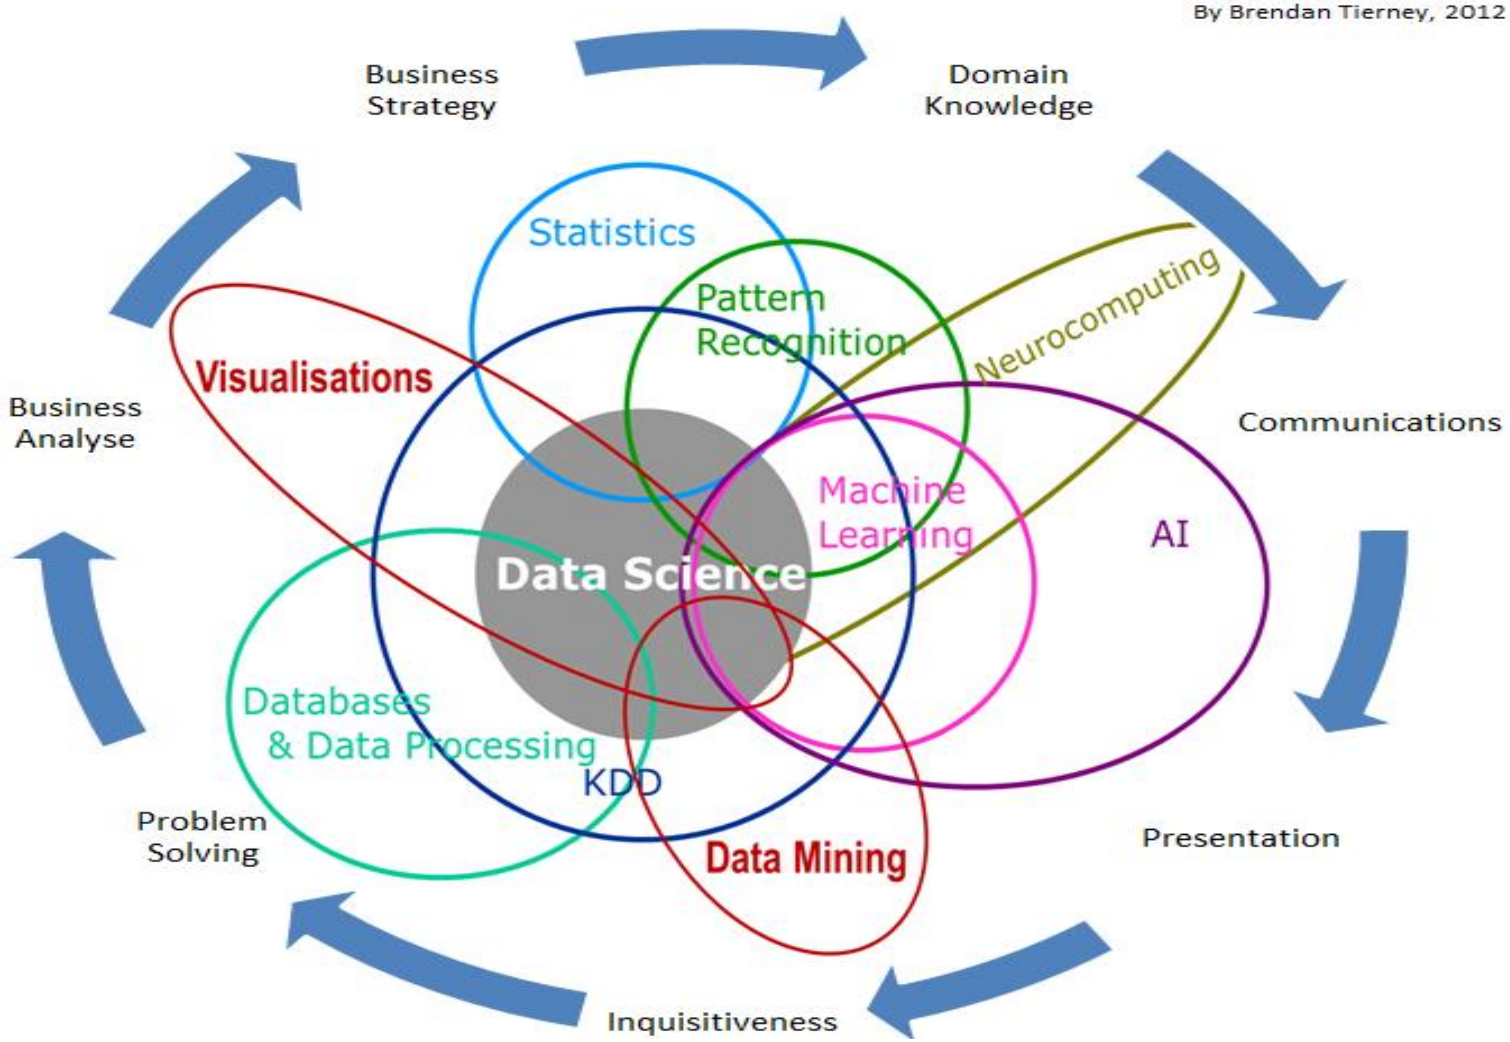

Data Science

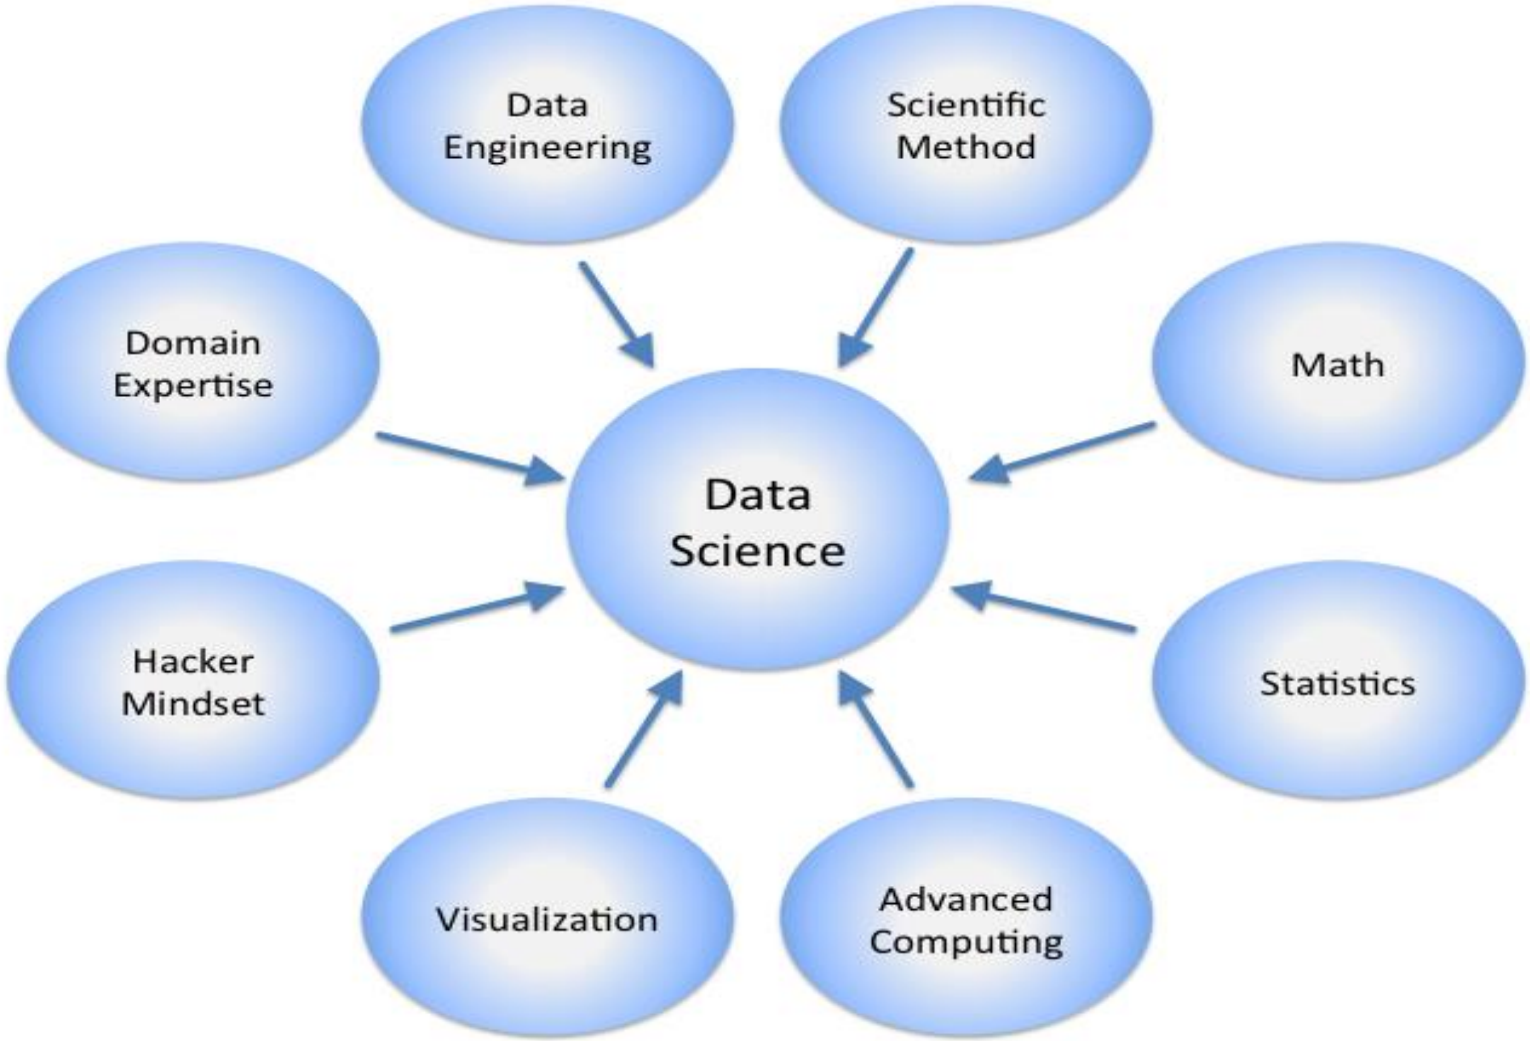

Why is it sexy?

- Gartner's

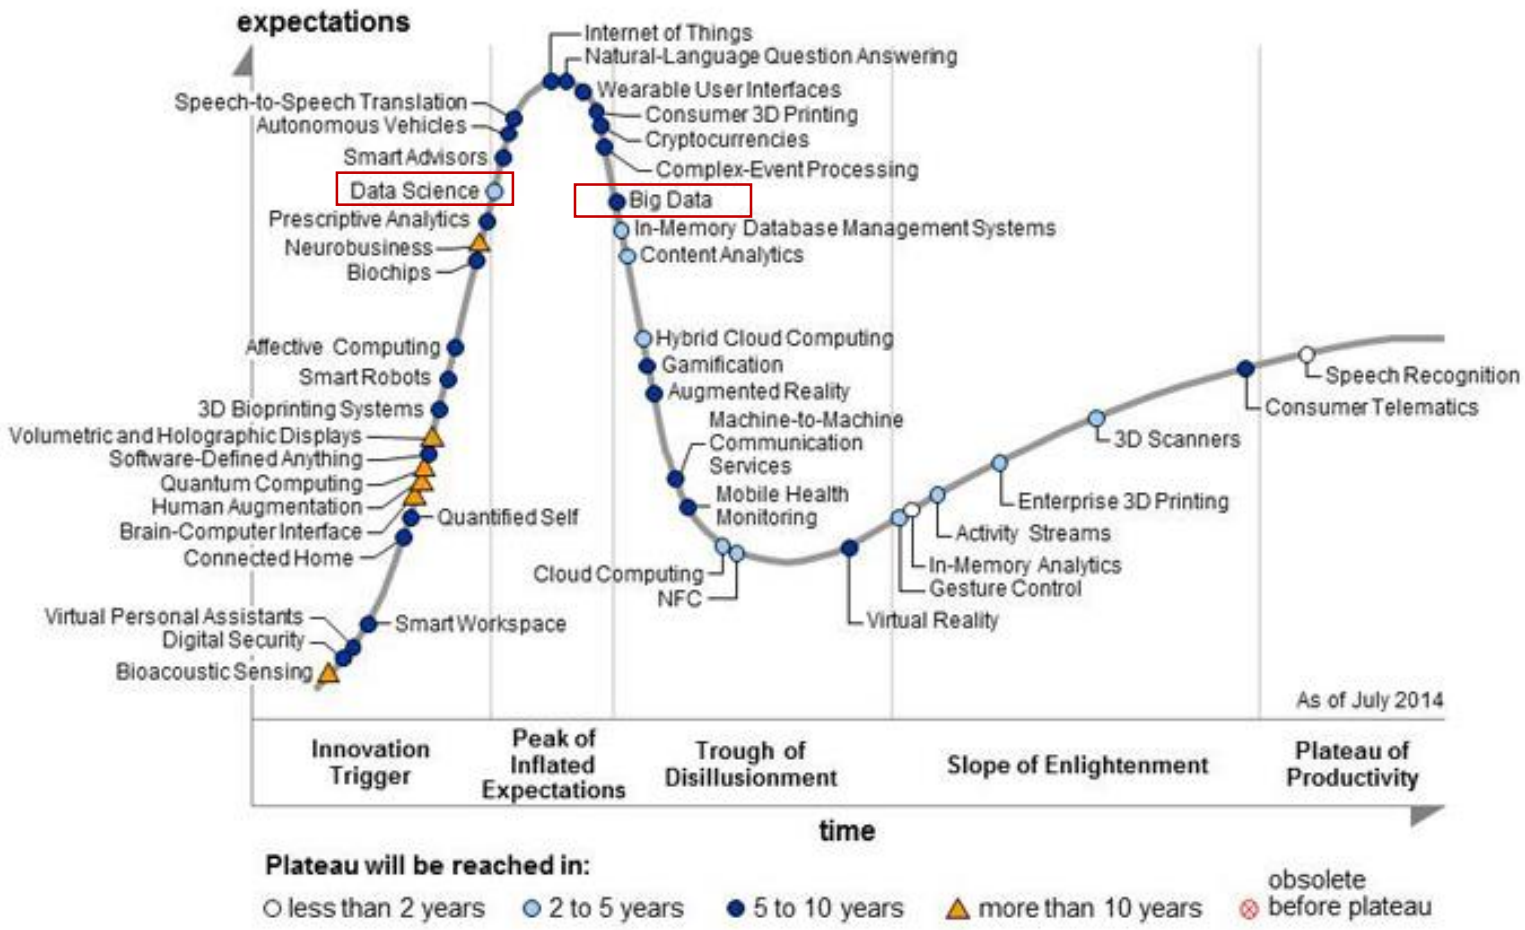

Data Scientists

- Data Scientist
  - The Sexiest Job of the 21st Century
- They find stories, extract knowledge. They are not reporters

Data Scientists

- Data scientists are the key to realizing the opportunities presented by big data. They bring structure to it, find compelling patterns in it, and advise executives on the implications for products, processes, and decisions

Concentration in Data Science

- Mathematics and Applied Mathematics
- Applied Statistics/Data Analysis
- Solid Programming Skills (R, Python, Julia, SQL)
- Data Mining
- Data Base Storage and Management
- Machine Learning and discovery

Machine Learning Problems

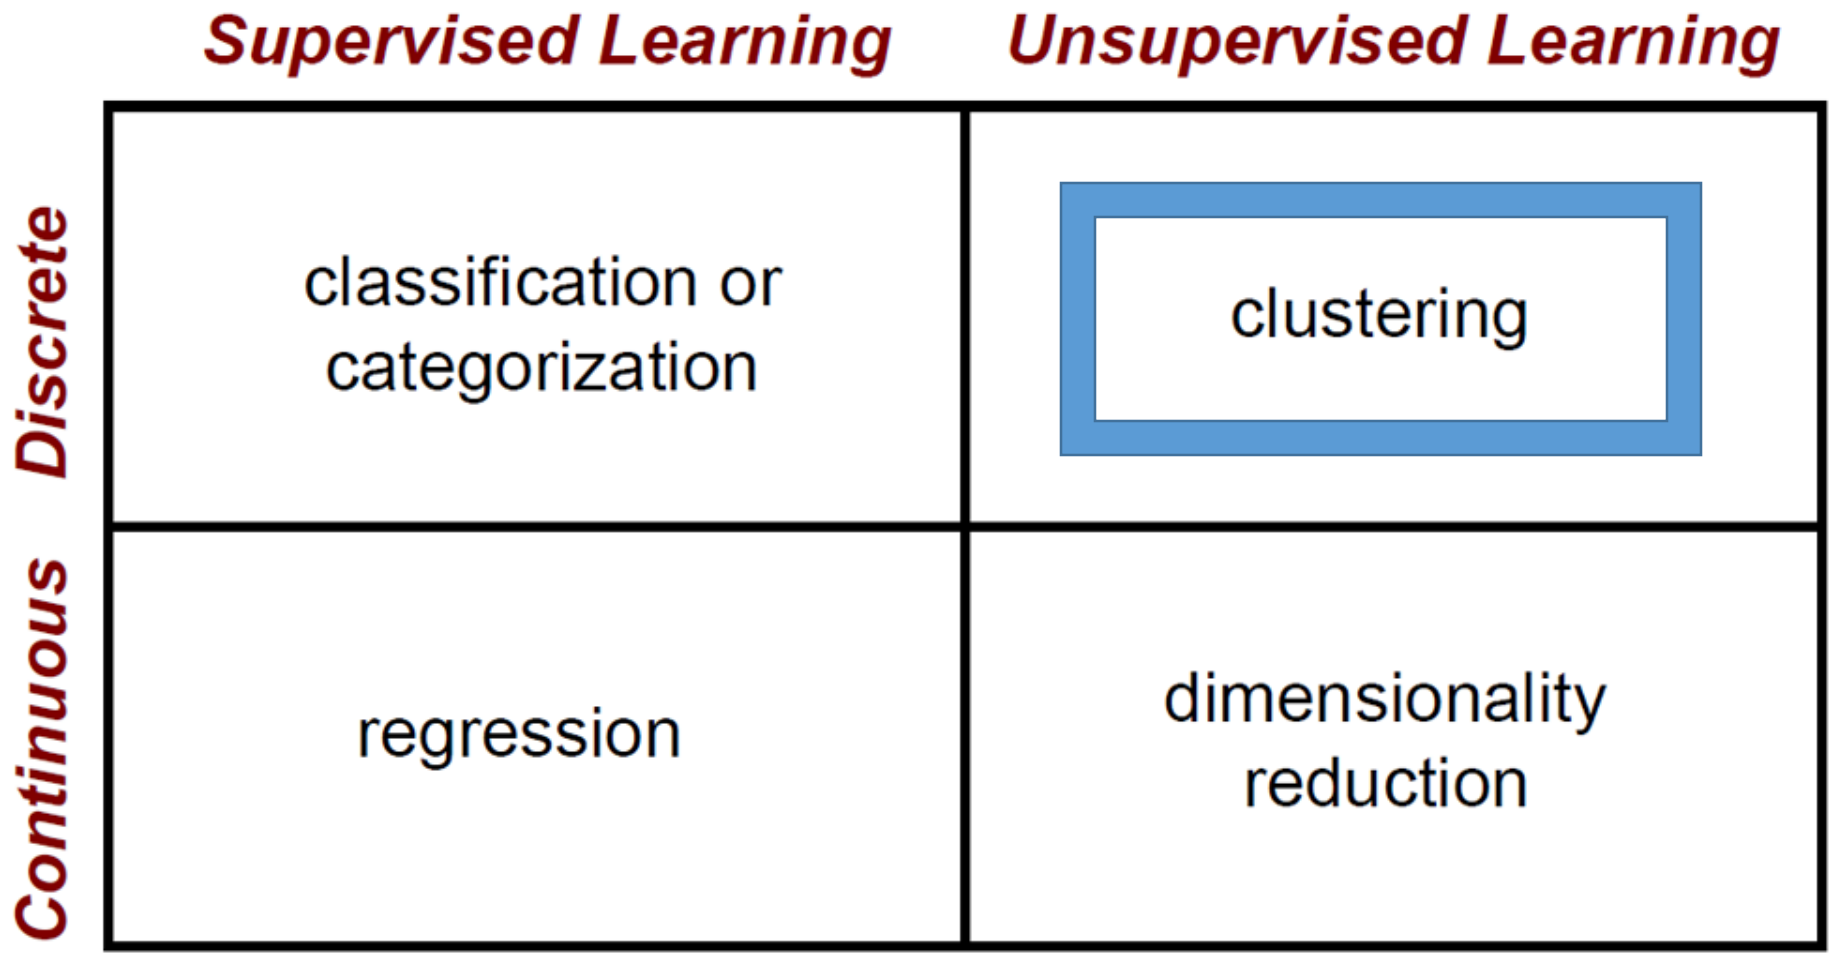

COMPLEMENTARIES AND DIFFERENCES

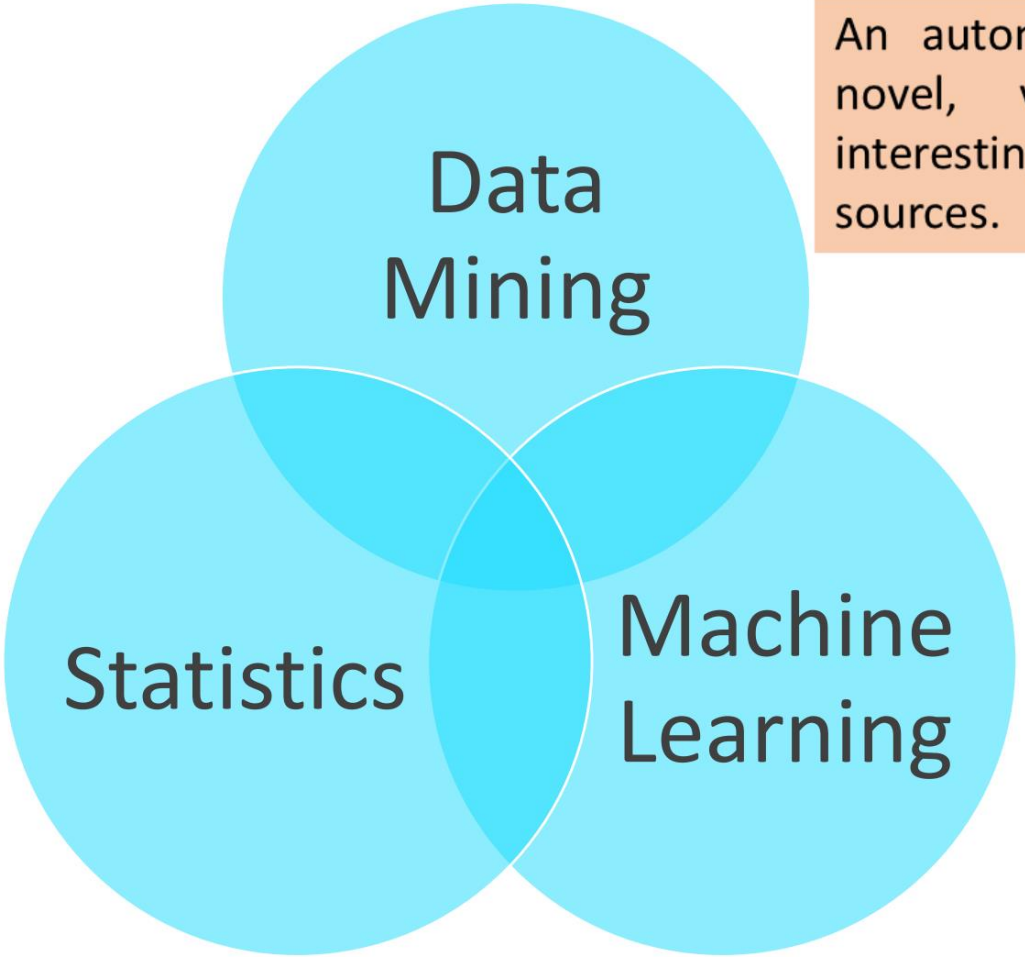

An automated process used to discover novel, valid, useful and potentially interesting knowledge from large data sources.

- Statistical analysis outputs: p-values, standard errors, regression models, principal components, discriminant score functions, ANOVA tables, control charts, descriptive statistics etc…
- translate statistical results into relevant information, careful formulation of findings is required

- Deals with representation and generalization
- Representation of data instances and functions evaluated on these instances
- Generalization is the property that the system will perform well on unseen instances

An Example of Machine Learning: Google's Self-Driving Car

COMPLEMENTARIES AND DIFFERENCES

Type of Method	Statistics	Machine Learning	Data Mining
Descriptive Methods	Statistical Methods	Rule format based	Mixed Types between Stat&ML
	p-values	Decision trees	
	Standard errors		
Clustering	Partitioning Clustering	Conceptual Clustering	All Types
	Hierarchical Clustering	Rule-Based Clustering	
Classification	Discriminat function	Neural nets	Mixed types
	K-NN classifier	Rule-based classifier	
	CART decision tree	Case-based reasoning	
		Decision tree induction	
Regression	Regression Methods	Regression Methods	Regression Methods
Association Rules		Association rule methods	Association rule methods
Visualizations	Dendrogram	Tree Representation Visualization	Tree Representation Visualization

THE DATA MINING CYCLE AND SOME TYPICAL APPLICATION DOMAINS

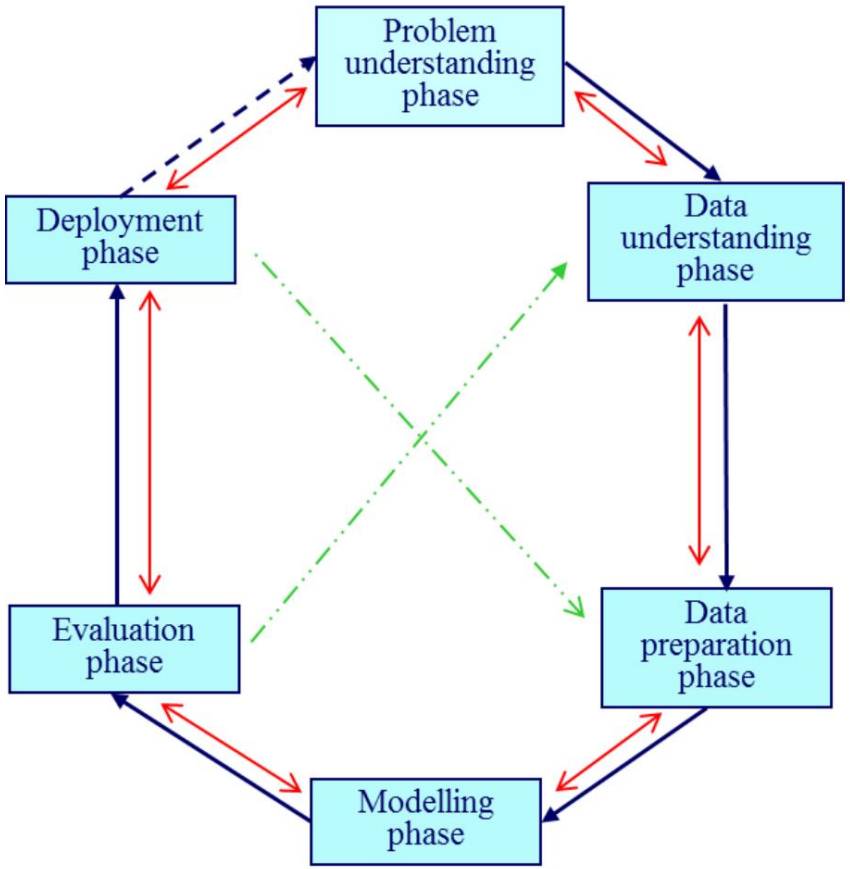

Database Marketing

Finance

Manufacturing

Health Insurance
Medicine/Patient records
Bioinformatics

Homeland Security

BioDefense

Text Mining

WWW mining
Crowdsourcing
Twitter

Optimization algorithms

- Example: $f = x_1 + x_2 + {x_1}^2 + x_1 e^{-x_2} + {x_2}^2 e^{-x_1}$

  Stationary points:

$$\nabla f = \mathbf{0} \Rightarrow \begin{cases} \dfrac{\partial f}{\partial x_1} = 1 + 2x_1 + e^{-x_2} - x_1 {x_2}^2 e^{-x_1} = 0 \\ \dfrac{\partial f}{\partial x_2} = 1 - x_1 x_2 e^{-x_2} + 2x_2 e^{-x_1} = 0 \end{cases}$$