University of Ljubljana, Faculty of Computer and Information Science

Intelligent Sytems

Prof Dr Marko Robnik-Šikonja

Ljubljana, October 2021

Lecturer

- Prof Dr Marko Robnik-Šikonja
- University of Ljubljana
  Faculty of Computer and Information Science
  Laboratory for Cognitive Modeling
- FRI, Večna pot 113, 2nd floor, to the right from the elevator
- marko.robnik@fri.uni-lj.si
- https://fri.uni-lj.si/en/employees/marko-robnik-sikonja
- (01) 4798 241
- Contact hour (see webpage)
  - currently, Wednesdays, 11:00 -12:00; for other times or video connection, email me
- **Research interests**: machine learning, artificial intelligence, natural language processing, network analytics, data science, data mining, algorithms and data structures
- **Teaching:** several courses from the area of data mining, algorithms, machine learning, and natural language processing
- **Software:** an author of three open source R packages from the area of predictive modelling and data analytics (CORElearn, semiArtificial, ExplainPrediction)

Assistants

- Blaž Škrlj, PhD student

- Tadej Škvorc, PhD student

- tutorials, assignments, work in R
  please, prepare questions!

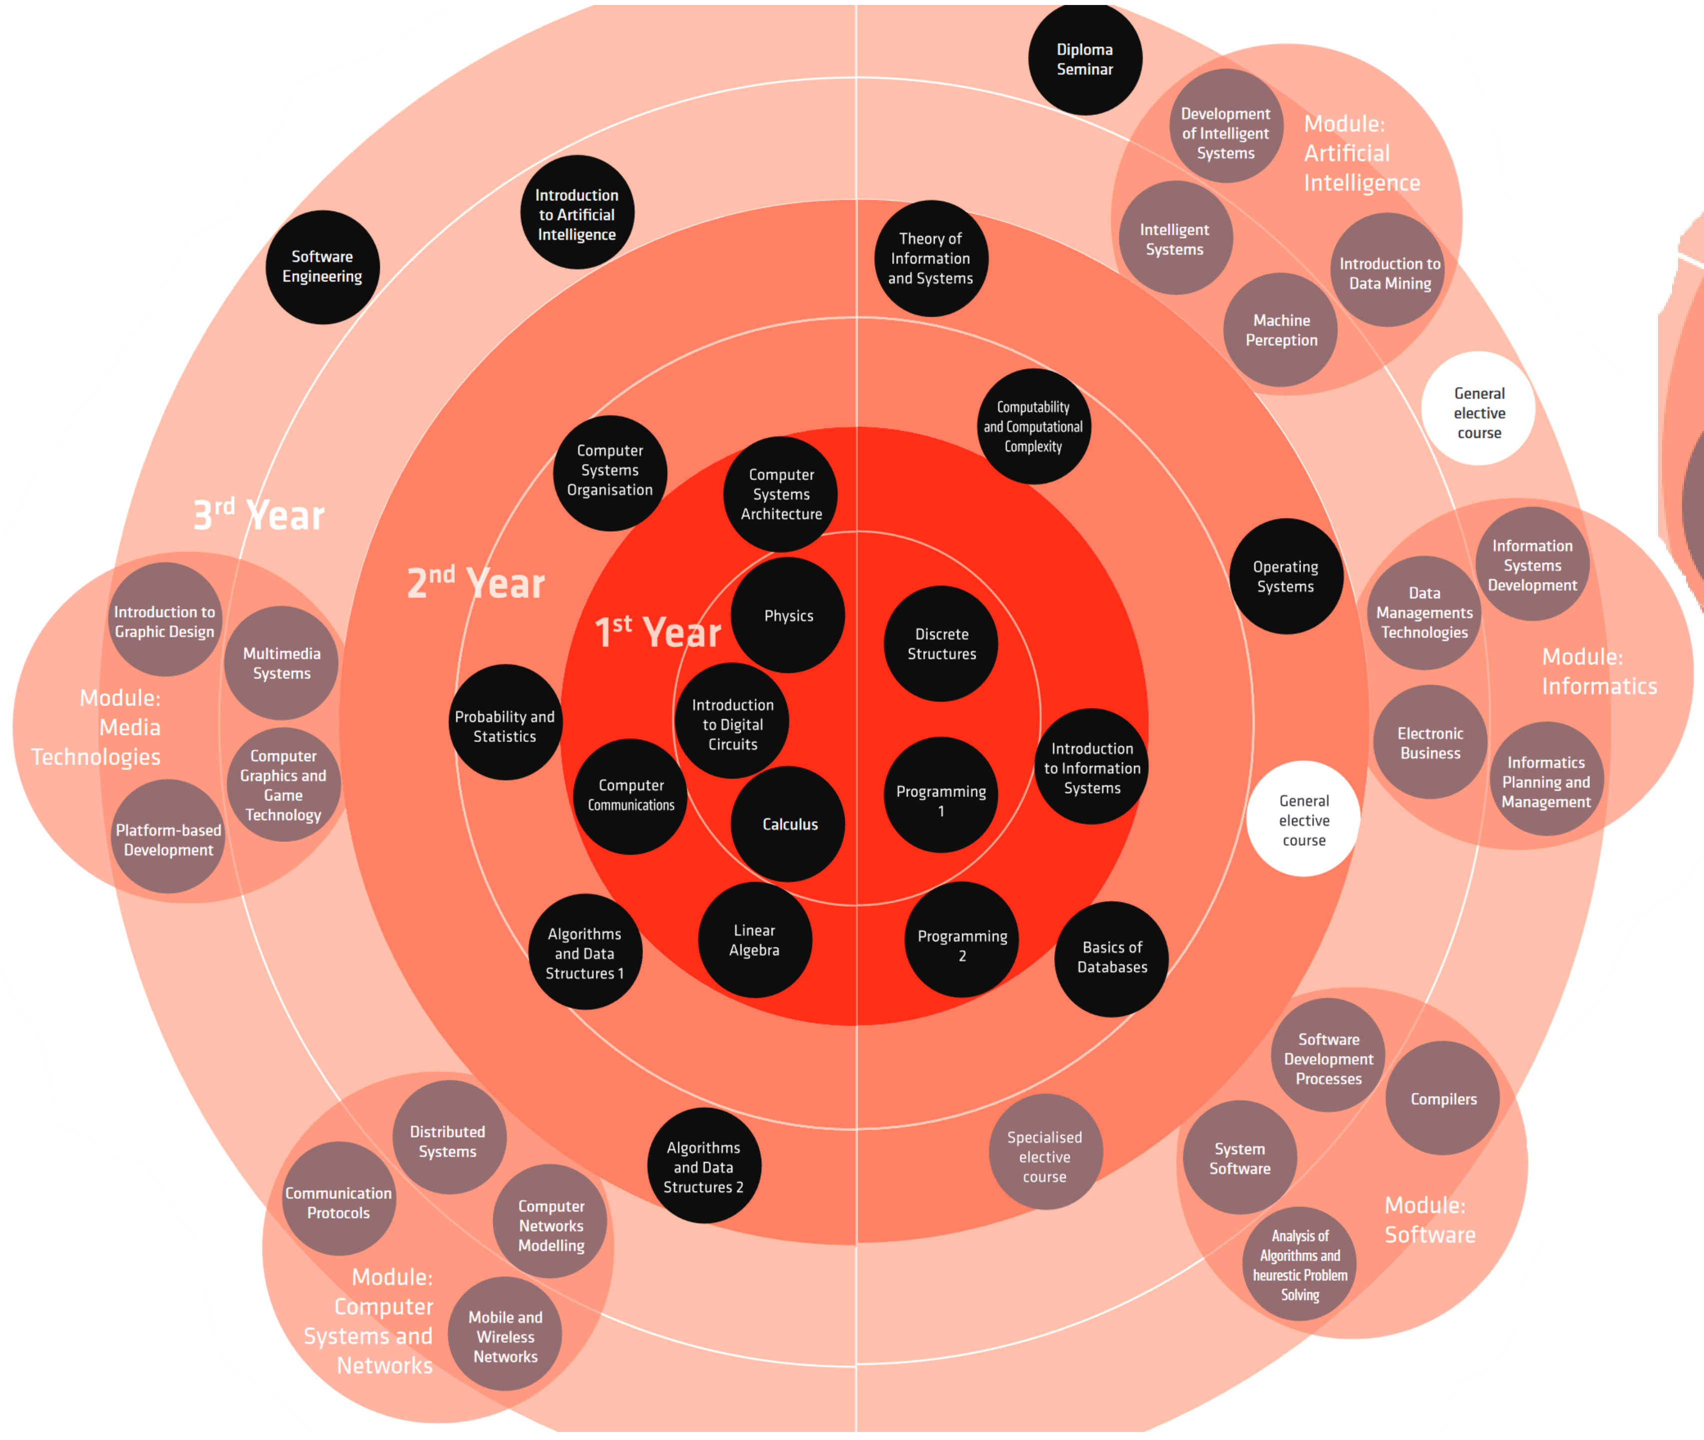

Diploma Seminar
Development of Intelligent Systems
Module: Artificial Intelligence
Introduction to Artificial Intelligence
Software Engineering
Theory of Information and Systems
Intelligent Systems
Introduction to Data Mining
Machine Perception
General elective course
Computability and Computational Complexity
Computer Systems Organisation
Computer Systems Architecture
3rd Year
2nd Year
1st Year
Operating Systems
Information Systems Development
Data Managements Technologies
Module: Informatics
Physics
Discrete Structures
Introduction to Graphic Design
Multimedia Systems
Module: Media Technologies
Probability and Statistics
Introduction to Digital Circuits
Introduction to Information Systems
Electronic Business
Informatics Planning and Management
Computer Graphics and Game Technology
Computer Communications
Programming 1
Calculus
General elective course
Platform-based Development
Algorithms and Data Structures 1
Linear Algebra
Programming 2
Basics of Databases
Software Development Processes
Compilers
Distributed Systems
Communication Protocols
Algorithms and Data Structures 2
Specialised elective course
System Software
Computer Networks Modelling
Module: Computer Systems and Networks
Module: Software
Analysis of Algorithms and heurestic Problem Solving
Mobile and Wireless Networks

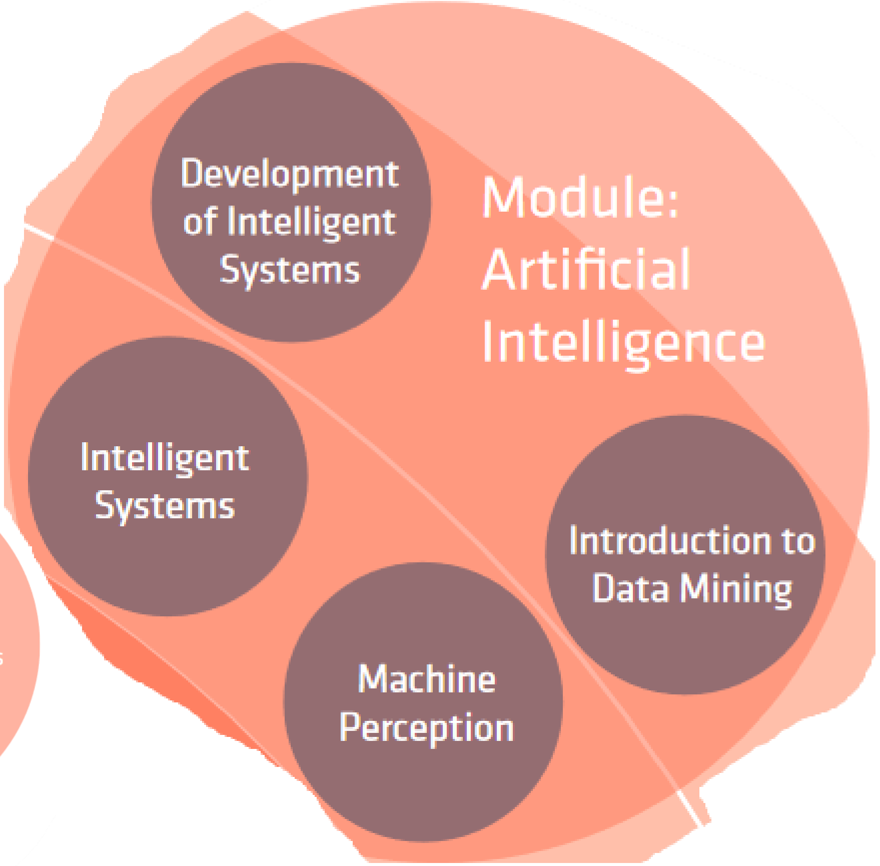

Development of Intelligent Systems
Module: Artificial Intelligence
Intelligent Systems
Introduction to Data Mining
Machine Perception

Syllabus

- nature inspired computing (genetic algorithms, genetic programming, ant colony optimization)
- basics of machine learning,
- similarity based learning, kNN
- decision rules and subgroup discovery
- ensemble methods
- support vector machines
- neural networks
- reinforcement learning
- natural language processing
- multiagent systems

Objectives

- students shall become acquainted with
  - nature inspired computing
  - machine learning
    - model selection and evaluation techniques
    - model comprehensibility and explanation
    - practical application of predictive modeling in R programming language and environment
  - natural language processing
  - reinforcement learning
  - multiagent systems
- practical use of theoretical knowledge on (almost) real-world problems
- awareness of domain expertise and ethical issues in data science
- increase the (mental) problem-solving toolbox with
  - predictive modeling techniques
  - reinforcement learning
  - result understanding, visualization and explanation approaches
- for a given prediction problem students shall be able to
  - transform it to a form suitable for predictive modeling
  - select and train an appropriate predictive model
  - evaluate the model and present the results in a comprehensible form and language.

Odd-one out

Concrete course goals]

Be able to explain

- difference between different types of machine learning models
- properties of models: bias, variance, generalization, hypothesis language
- properties of the following models: kNN, decision rules, bagging, boosting, random forests, stacking, SVM, neural networks
- properties and purpose of evaluation approaches and metrics: cross-validation, bootstrapping, ROC curves, sensitivity, specificity etc.
- inference methods for predictive methods and explanation of predictions
- when and why to apply reinfocement learning
- how to prepare and process text in a text mining
- when and how and to optimize a problem using genetic algorithms

Build and evaluate models in R

- visualize the data set and created models
- prepare data into a suitable form suitable for modeling algorithms
- apply classification and regression models to solve a prediction task with a given data set
- estimate error of models using statistically valid approaches
- select models and tune their parameters using cross-validation and bootstrapping
- visualize models and explain their predictions
- given a new data set select appropriate modeling technique and evaluate the created model

Why R?

- [IEEE Spectrum 2019 list of popular programming languages]

Rank	Language	Type	Score
1	Python		100.0
2	Java		96.3
3	C		94.4
4	C++		87.5
5	R		81.5
6	JavaScript		79.4
7	C#		74.5
8	Matlab		70.6
9	Swift		69.1
10	Go		68.0

Nature inspired computing

- genetic algorithms
- genetic programming
- ant colony optimization

Introduction to statistical predictive modelling

- Learning as modelling: data, evidence, background knowledge, predictive models, hypotheses, learning as optimization, learning as search, criteria of success, inductive learning, generalization.
- Classification and regression: supervised and unsupervised learning, learning discrete and numeric functions, learning relations, learning associations.
- Simple classification models: nearest neighbor, decision rules

Model selection

- Bias and variance: error decomposition, trade-off, estimating bias and variance
- Generalization performance: training and testing set error, cross-validation, evaluation set, bootstrapping.
- Performance measures: confusion matrix, sensitivity and specificity, ROC curves, AUC, cost-based classification.
- Parameter tuning: regularization, MDL principle.
- Calibration of probabilities: binning, isotonic regression.
- No free lunch theorem.

Ensemble methods

- Model averaging, why ensembles work.
- Tree based ensembles: bagging, boosting, random forests.
- MARS and AODE ensembles.
- Stacking.

Kernel methods

- SVM for classification and regression: kernels, support vectors, hyperplanes.
- SVM for more than two classes: one vs. one, one vs. all.

Neural networks

- perceptron,
- backpropagation,
- RBF networks,
- setting structure of networks
- deep neural networks
- autoencoders
- GANs
- the role of embeddings

Explaining prediction models

- Model comprehensibility, visualization and knowledge discovery.
- General methodology for explaining predictive models.
- Model level and instance level explanations, methods SHAP, LIME, EXPLAIN, and IME.

Learning with special settings

- imbalanced data,
- multi-task learning,
- multi-label learning.

Reinforcement learning

- basics
- Markov decision problem
- Q learning

Natural language processing

- text preprocesing
- text representation
- text similarity
- text classification
- sentiment analysis

Multiagent systems

- types of agents
- agent architectures
- distributed constraint satisfaction
- distributed path finding

Obligations

- 5 quizzes
- three projects, 50 points
- written exam, 50 points

Grading

Obligation	% of total	subject to
Five quizzes	0%	≥ 50% alltogether
Three projects	50%	≥ 50% each
Written exam	50%	≥ 50%

Learning materials

- learning materials in the eClassroom
- slides
- links to textbook and papers
- R code and examples
- links to data sets
- install the open-source systems R and RStudio

Readings

- James, G., Witten, D., Hastie, T., & Tibshirani, R. (2013).
  An introduction to statistical learning with applications in R. Springer, New York.
  - freely available from authors' homepages
  - also code and slides from authors and Abbass Al Sharif (some used in this course)

Further readings:

- Friedman, J., Hastie, T., & Tibshirani, R. (2001). *The elements of statistical learning*. Springer, Berlin.
- Ian Goodfellow, Yoshua Bengio, and Aaron Courville (2016). *Deep Learning*. MIT press, freely available from authors' homepages
- François Chollet (2017). Deep Learning with Python. Manning.
- Kononenko, I., Robnik-Šikonja, M.: Inteligentni sistemi. Založba FE in FRI, 2010 (in Slovene)
- scientific papers
- many excellent machine learning and data mining courses on Coursera and edX

BTW: retention of learning

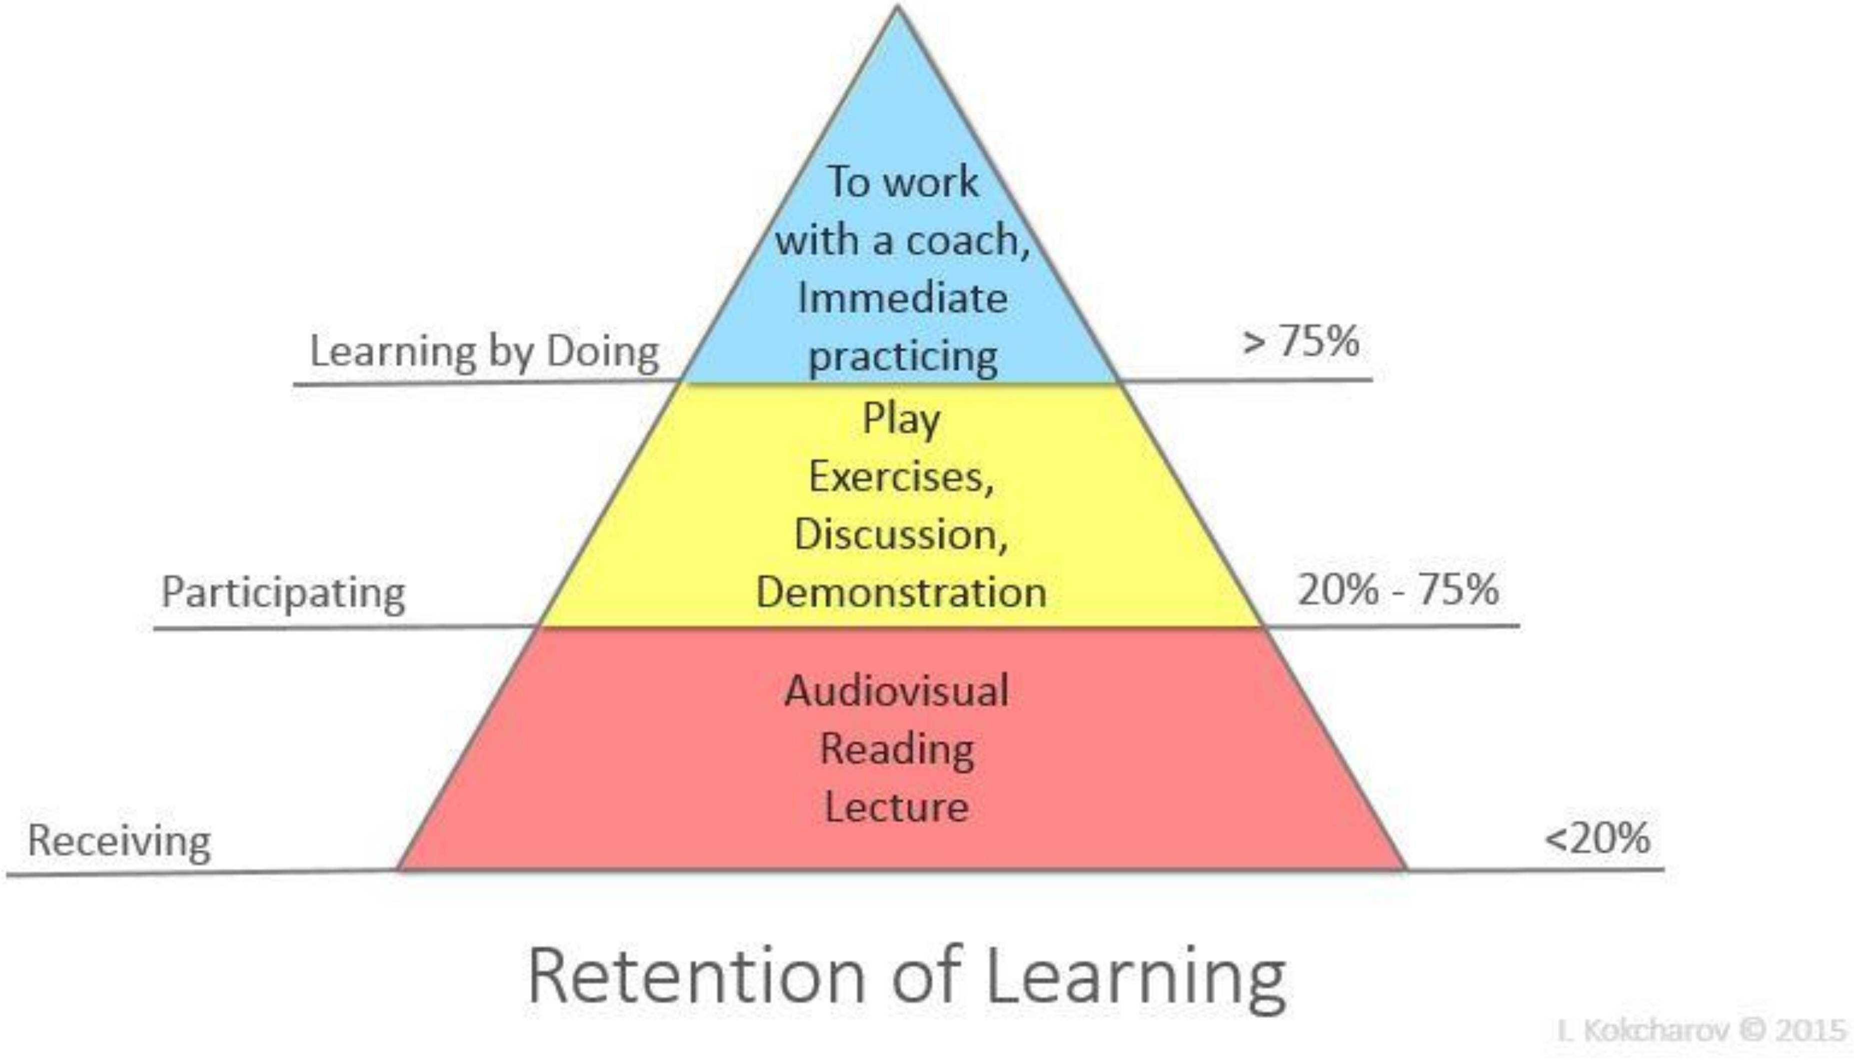

Data Science

- good job perspective
- Forbes list of the most promising jobs in the USA

Top 10 Best Jobs in America in 2020

Rank	Job Title	Median Base Salary	Job Satisfaction	Job Openings
1	Front End Engineer	$105,240	3.9	13,122
2	Java Developer	$83,589	3.9	16,136
3	Data Scientist	$107,801	4.0	6,542
4	Product Manager	$117,713	3.8	12,173
5	Devops Engineer	$107,310	3.9	6,603
6	Data Engineer	$102,472	3.9	6,941
7	Software Engineer	$105,563	3.6	50,438
8	Speech Language Pathologist	$71,867	3.8	29,167
9	Strategy Manager	$133,067	4.3	3,515
10	Business Development Manager	$78,480	4.0	6,560

Source: Glassdoor Economic Research (Glassdoor.com/research)

- Thomas H. Davenport, D.J. Patil: Data Scientist: The Sexiest Job of the 21st Century. *Harvard Business Review*, October 2012

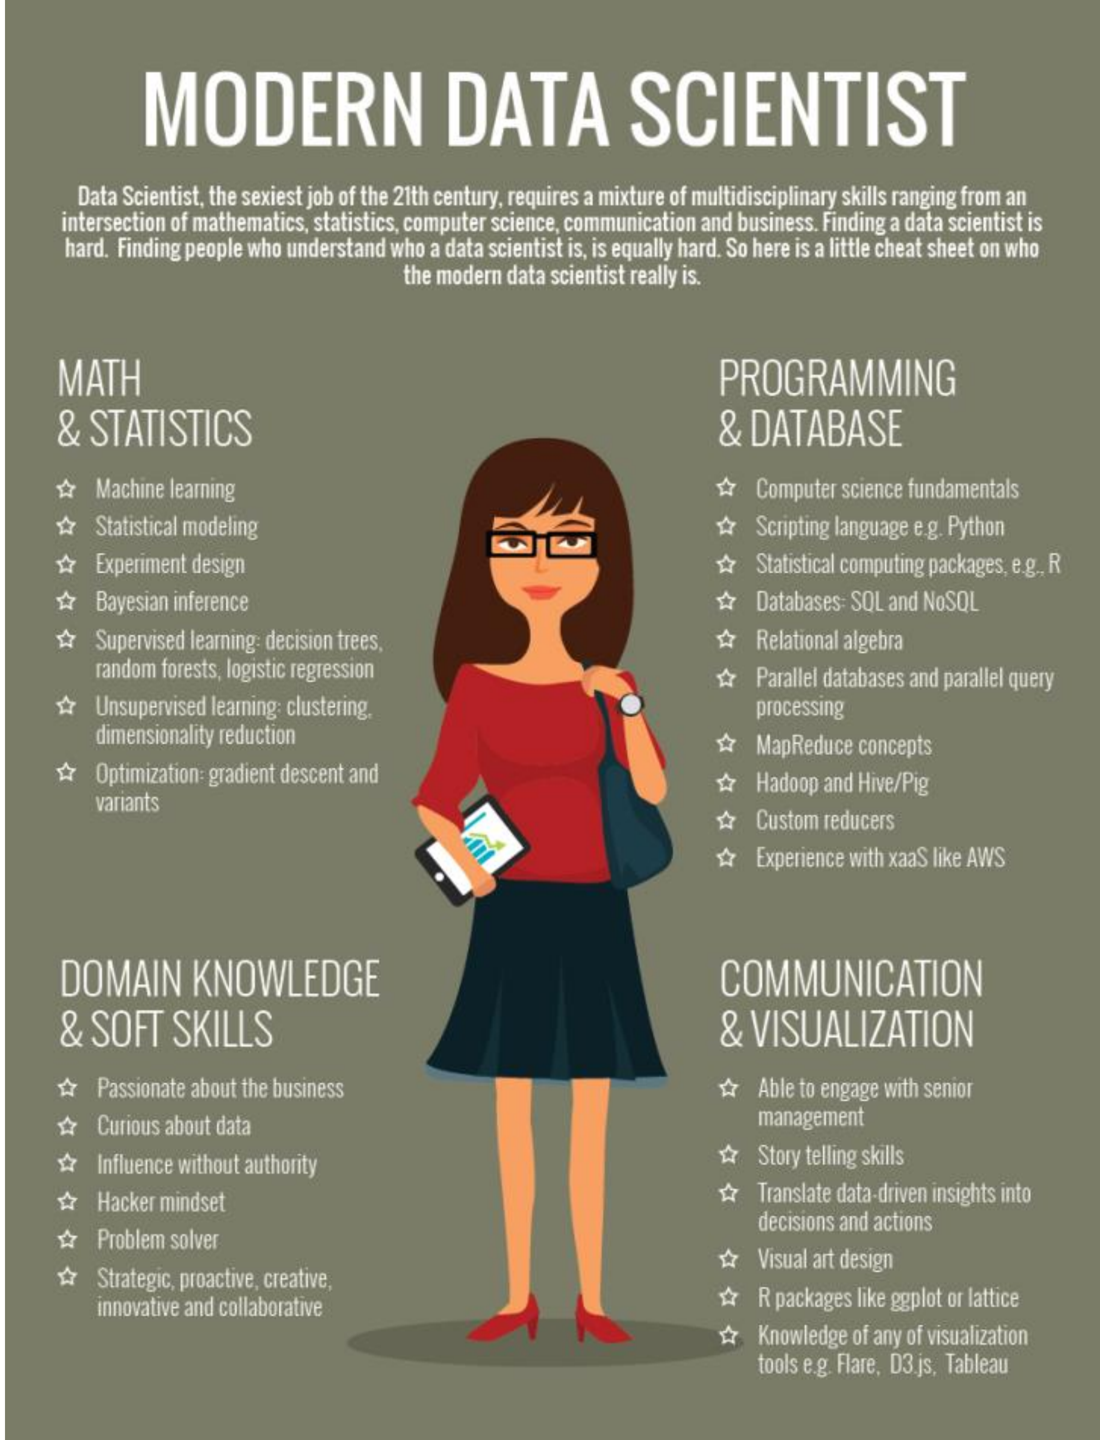


I THOUGHT I HAD THE SEXIEST JOB OF THE 21ST CENTURY!
KDNUGGETS.COM

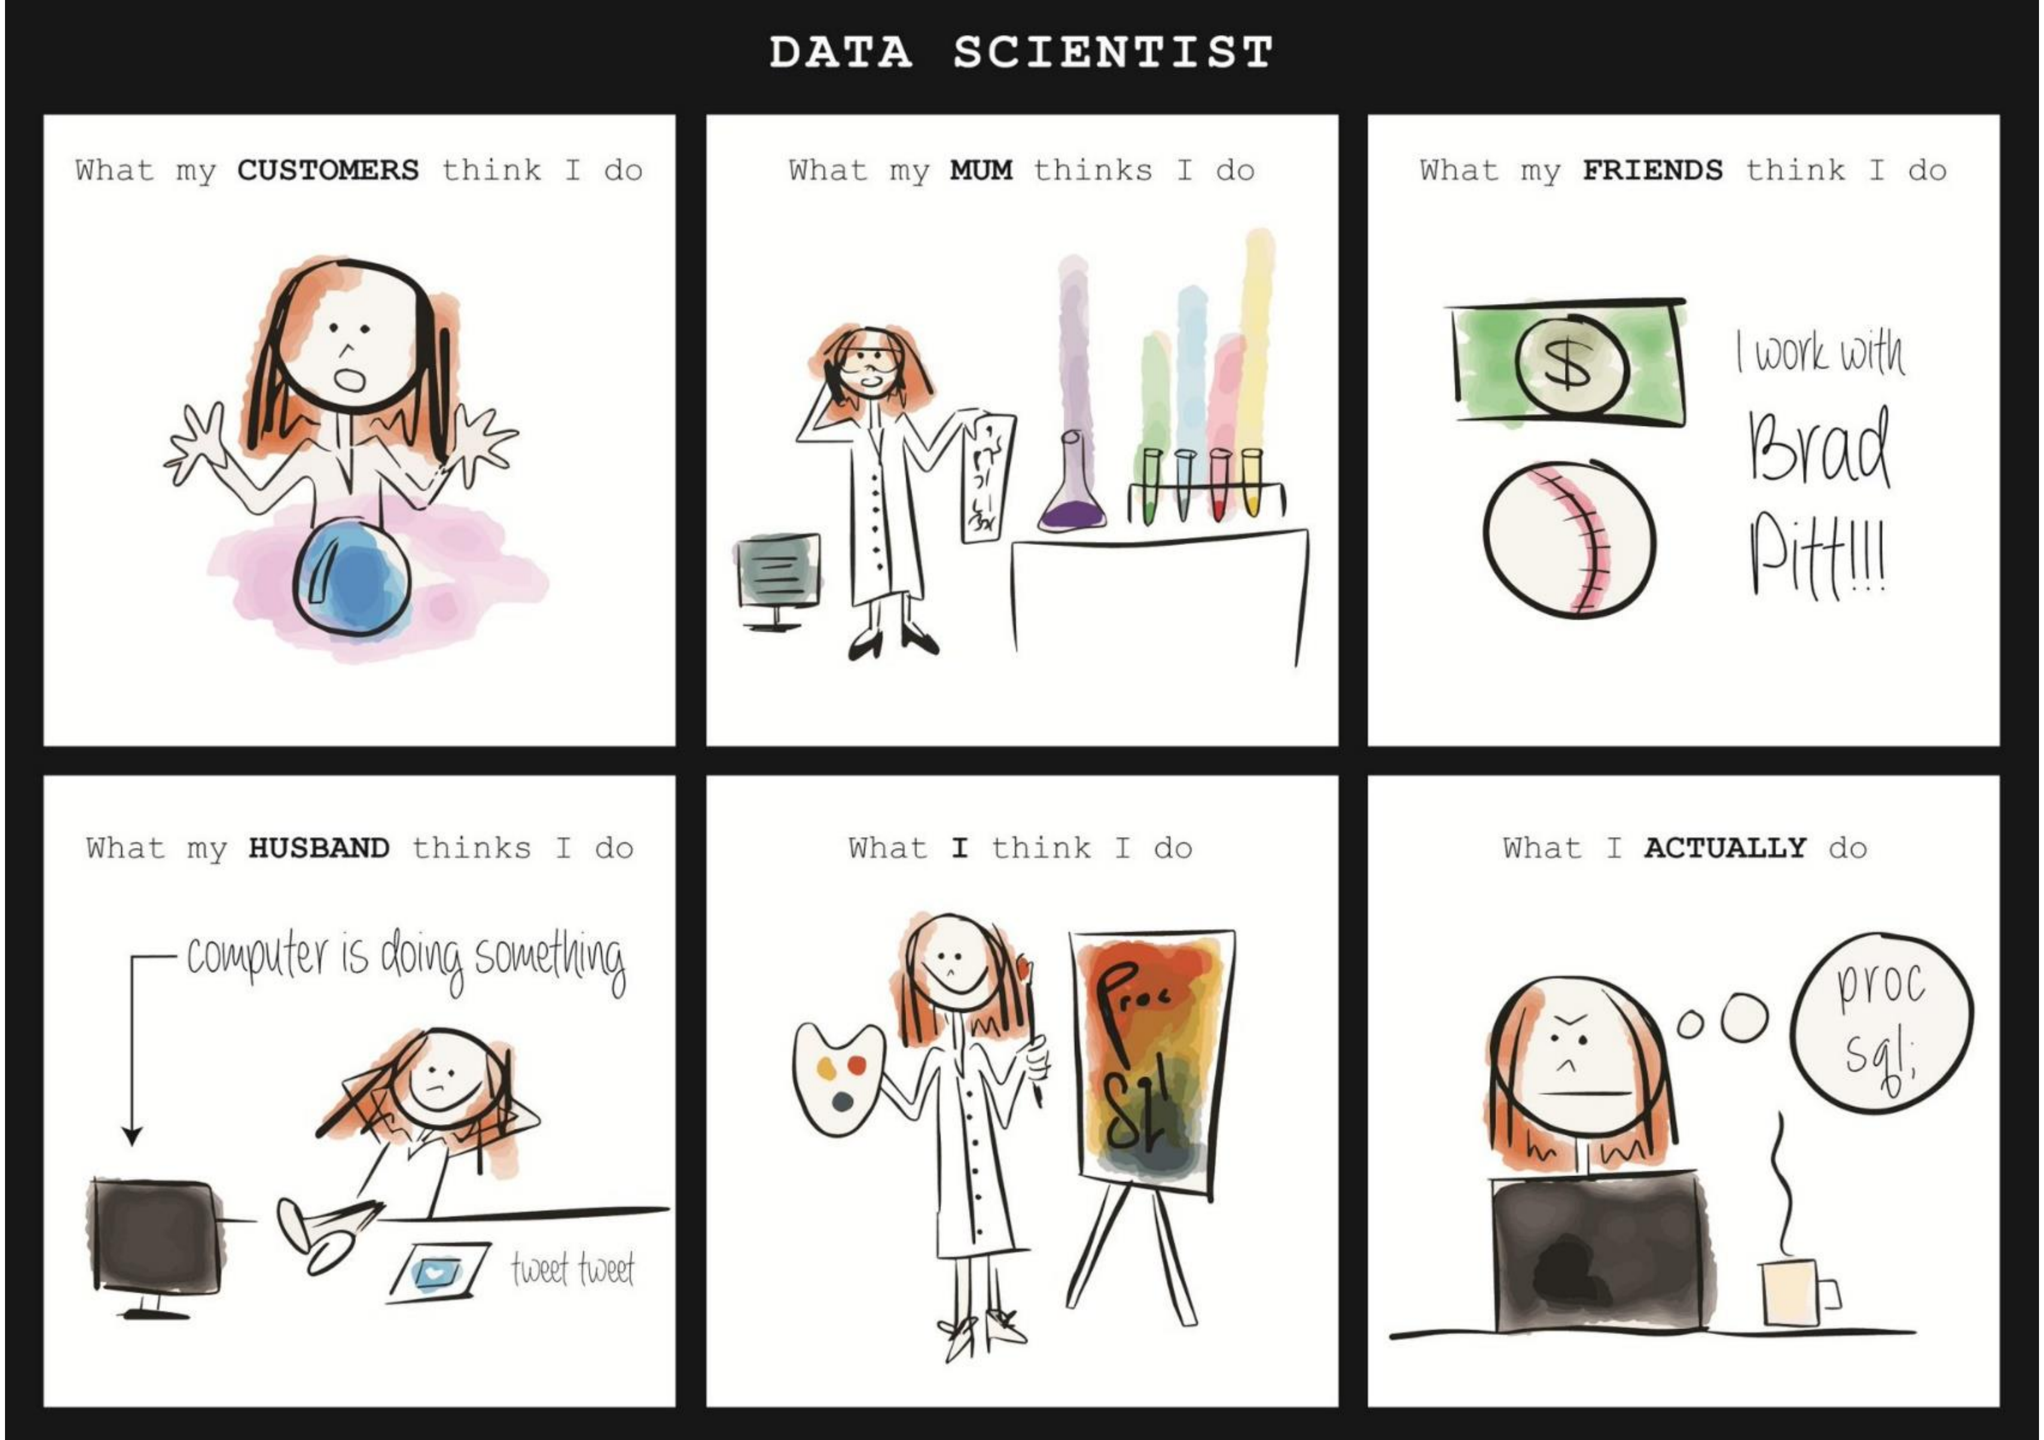
DATA SCIENTIST
What my CUSTOMERS think I do
What my MUM thinks I do
What my FRIENDS think I do
I work with Brad Pitt!!!
What my HUSBAND thinks I do
computer is doing something
tweet tweet
What I think I do
Proc
Sql
What I ACTUALLY do
proc sql;

Intelligent systems and media

Will robots destroy us?
Will they take our jobs?
Will we still need a driving licence?
Will we still need doctors?
How will humanoid robots evolve?
What about cyborgs?
What is artificial general intelligence?
What is technological singularity?

New prophets of tehnological singularity

Elon Musk says humans must become cyborgs to stay relevant. Is he right?

Sophisticated artificial intelligence will make 'house cats' of humans, claims the entrepreneur, but his grand vision for mind-controlled tech may be a long way off

Some scientific responses

- Rodney Brooks: The Seven Deadly Sins of Predicting the Future of AI. https://rodneybrooks.com/the-seven-deadly-sins-of-predicting-the-future-of-ai/ tudi MIT Technology Review
- Marko Robnik-Šikonja: Is artificial intelligence a (job) killer?. The Conversation, Jul. 2017 https://theconversation.com/is-artificial-intelligence-a-job-killer-80473
- …

Short history of optimism

- starting in 1950s,
  1956 Dartmouth conference
- great expectations, enormous underestimation of problem difficculty
- AI winter (2 x)

1958, H. A. Simon and Allen Newell: “… within ten years a digital computer will discover and prove an important new mathematical theorem.”

1965, H. A. Simon: “… machines will be capable, within twenty years, of doing any work a man can do.”

1967, Marvin Minsky: "Within a generation ... the problem of creating 'artificial intelligence' will substantially be solved.”

1970, Marvin Minsky: "In from three to eight years we will have a machine with the general intelligence of an average human being."

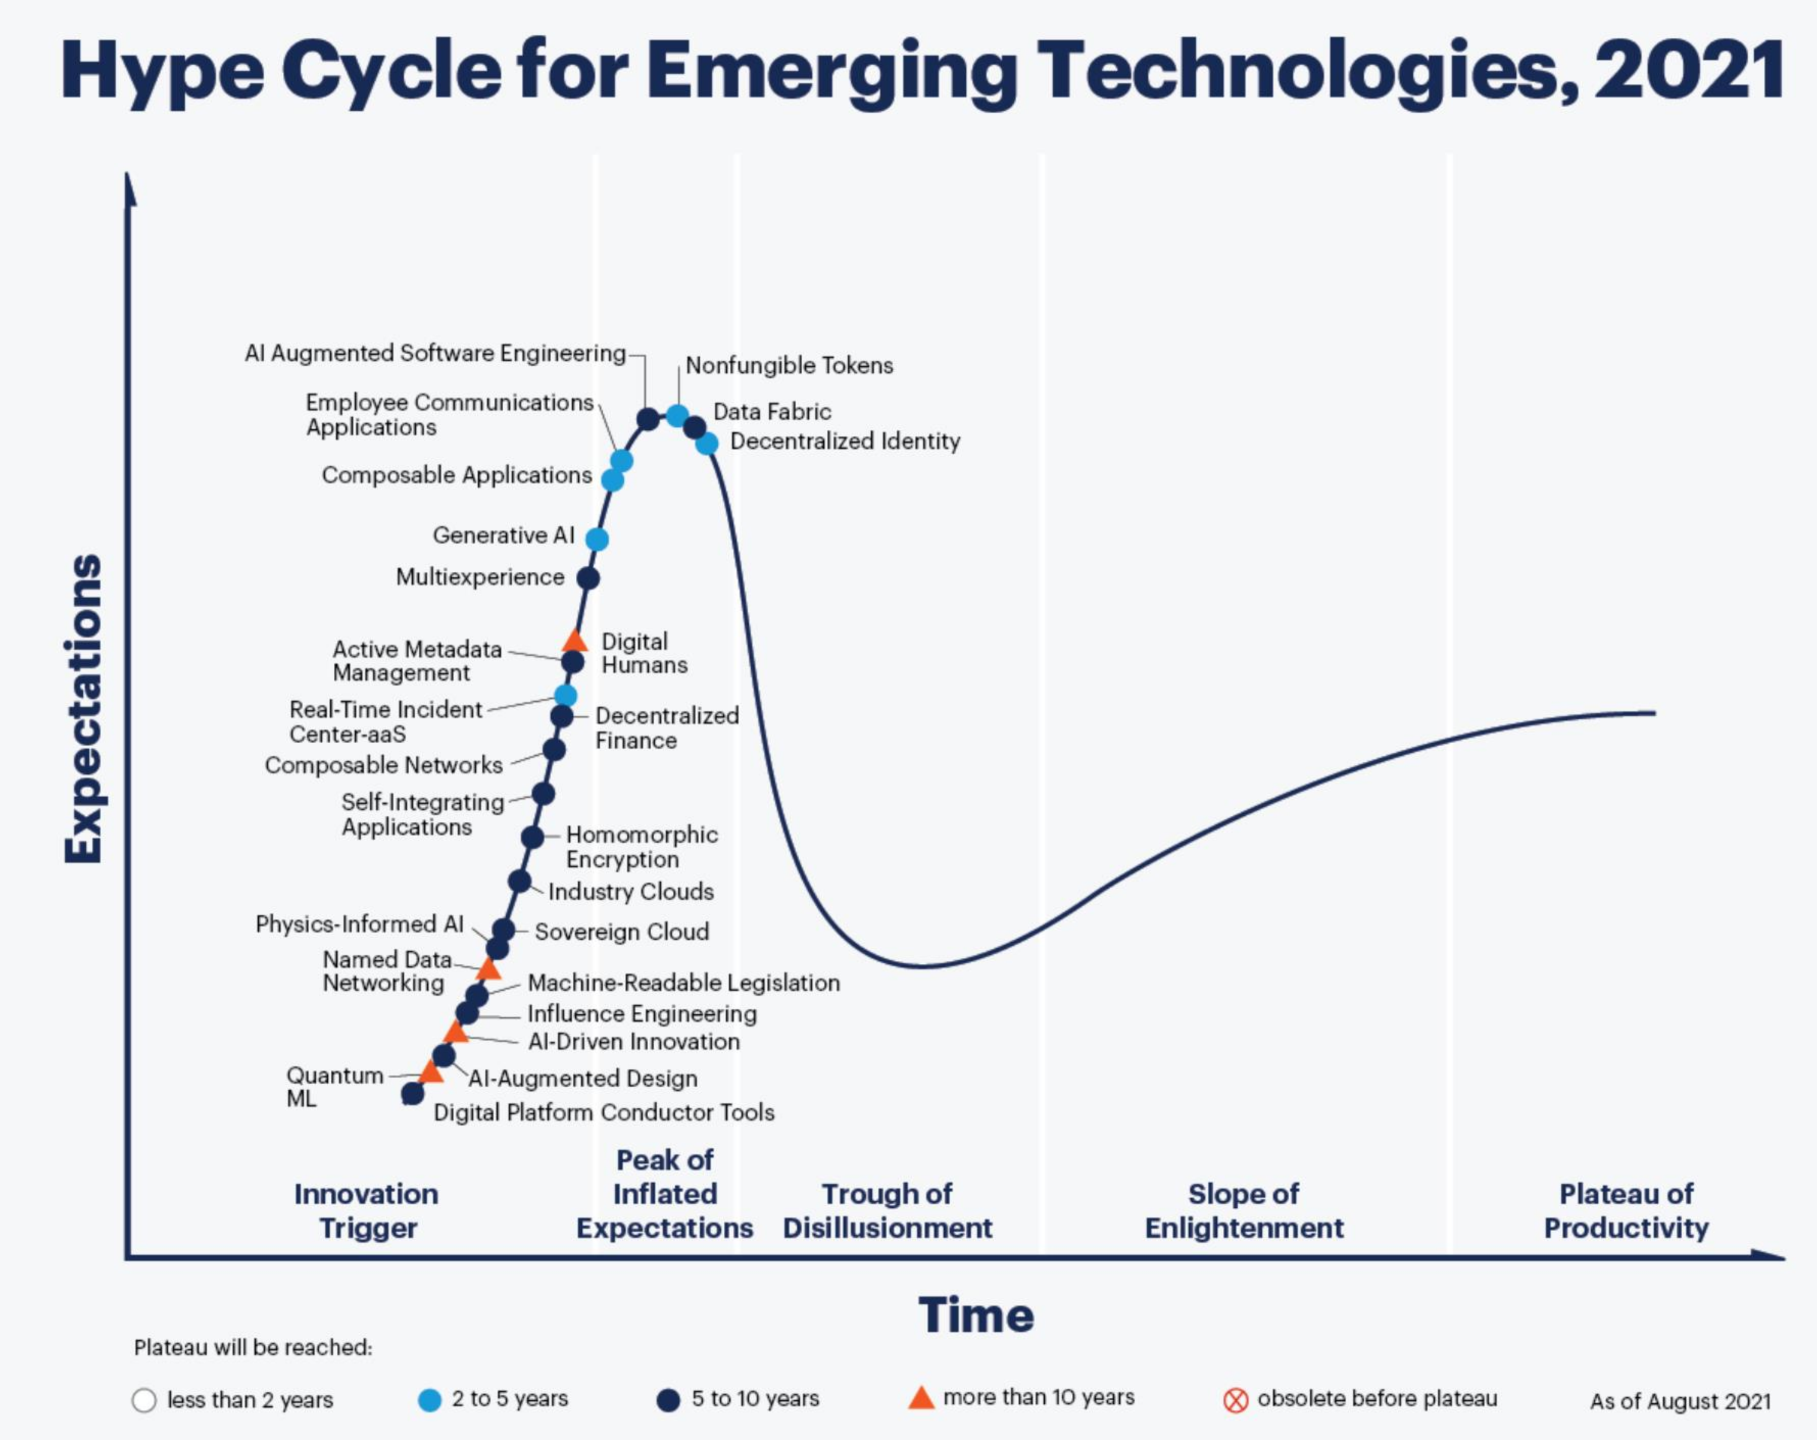
Hype Cycle for Emerging Technologies, 2021
Expectations
AI Augmented Software Engineering
Nonfungible Tokens
Employee Communications Applications
Data Fabric
Decentralized Identity
Composable Applications
Generative AI
Multiexperience
Active Metadata Management
Digital Humans
Real-Time Incident Center-aaS
Decentralized Finance
Composable Networks
Self-Integrating Applications
Homomorphic Encryption
Industry Clouds
Physics-Informed AI
Sovereign Cloud
Named Data Networking
Machine-Readable Legislation
Influence Engineering
AI-Driven Innovation
Quantum ML
AI-Augmented Design
Digital Platform Conductor Tools
Innovation Trigger
Peak of Inflated Expectations
Trough of Disillusionment
Slope of Enlightenment
Plateau of Productivity
Time
Plateau will be reached:
less than 2 years
2 to 5 years
5 to 10 years
more than 10 years
obsolete before plateau
As of August 2021

Hype Cycle for Emerging Technologies, 2020

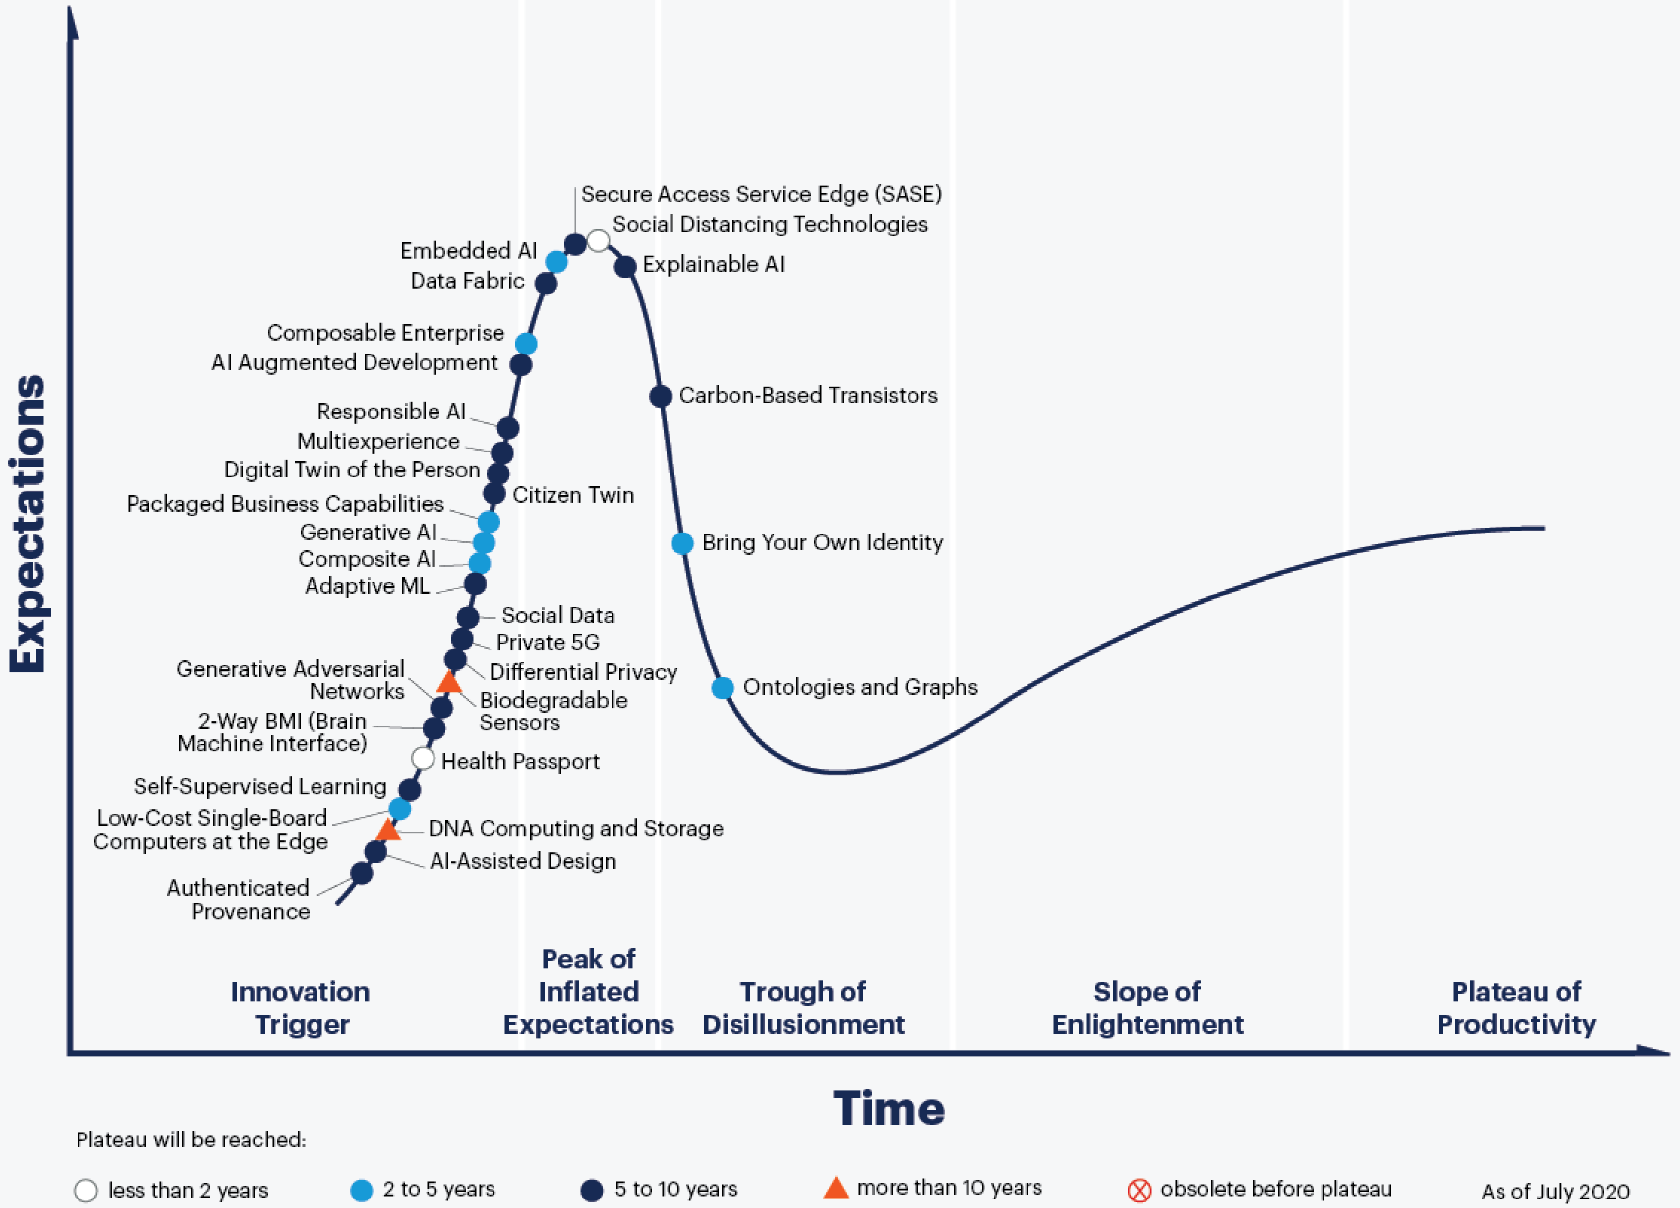

Gartner Hype Cycle for Emerging Technologies, 2019

Hype Cycle for Artificial Intelligence, 2021

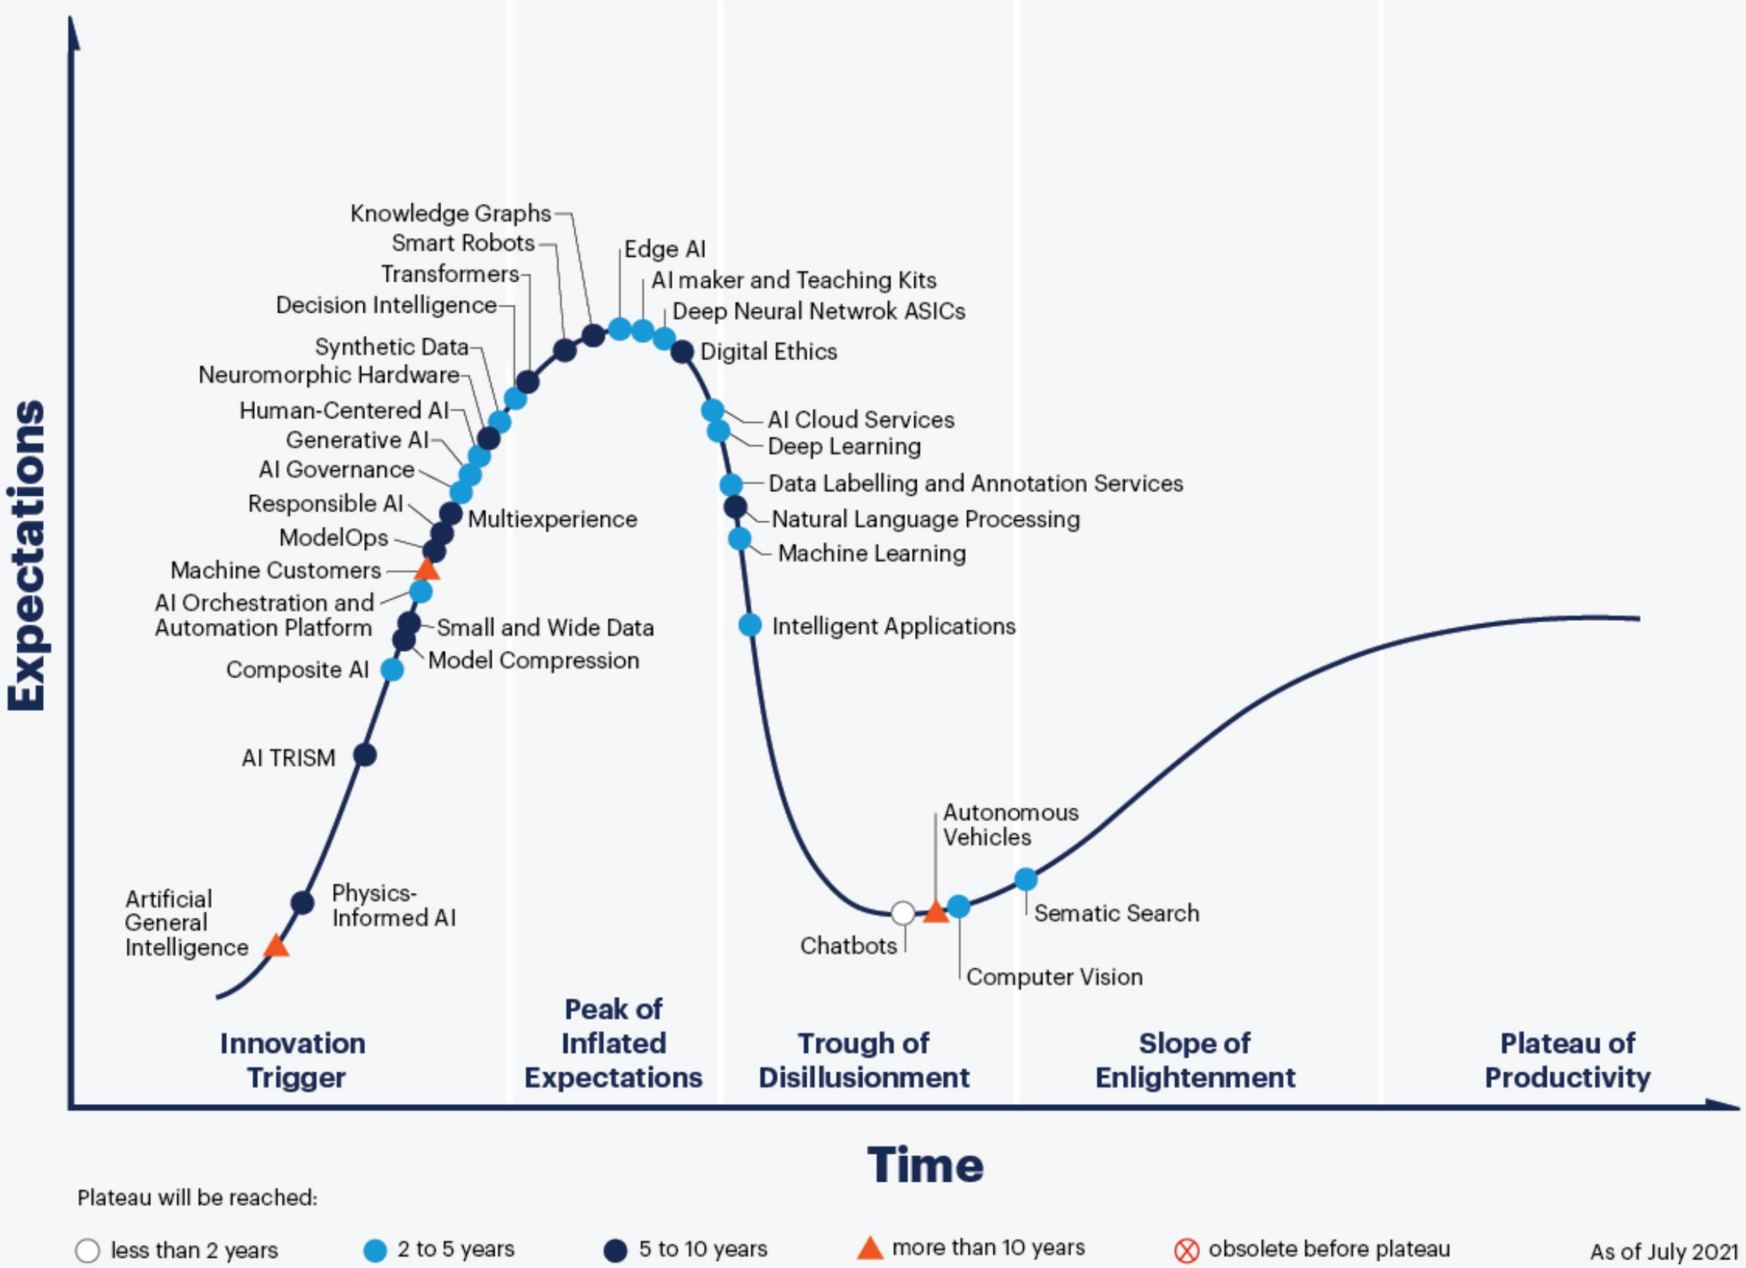

Hype Cycle for Artificial Intelligence, 2020