1. Εισαγωγή

Τα τελευταία χρόνια επιτελείται τεράστια ανάπτυξη σε τεχνολογίες και εφαρμογές σχετικές με :

- Τεχνητή Νοημοσύνη και Μηχανική Μάθηση
- Αναπαράσταση και Διαχείριση Γνώσης
- Αναλυτική Μεγάλων Δεδομένων (Δομημένων ή Μη)
- Εφαρμογές στην βιομηχανία (IoT), στην υγεία, στις υπηρεσίες του Δημοσίου, στην ανάπτυξη έξυπνων σπιτιών, πόλεων, στα κοινωνικά δίκτυα, κά.

Εκπαίδευση - Υποδομές

- Τα Πανεπιστήμια απαντούν στις νέες προκλήσεις εμπλουτίζοντας τα προγράμματα σπουδών ώστε να δημιουργήσουν εξειδικευμένα στελέχη και επιστήμονες που θα εκμεταλλευτούν τους τεράστιους όγκους των δεδομένων
- Οι Εταιρείες αναπτύσσουν τεχνολογίες που μπορούν να αναλύουν μεγάλα δεδομένα, δημιουργούν εξειδικευμένες βάσεις δεδομένων, σχεδιάζουν και υλοποιούν μεθόδους συλλογής, επεξεργασίας των δεδομένων και εξαγωγής συμπερασμάτων που κάνουν τις παρεχόμενες υπηρεσίες πιο αποτελεσματικές.

Παραγωγή Πολύπλοκων Δεδομένων

- **Τεράστιοι Ογκοι Δεδομένων**: Peta (1000.000.000.000.000) Χαρακτήρες
- **Διαφορετικές Μορφές**: Κείμενα, Εικόνες, Ηχος, Βίντεο, Οντολογίες, Γράφοι, Πίνακες
- **Κατανενημένα, Δυναμικά, Θορυβώδη Δεδομένα**
- **Διασυνδεδεμένα Ανοικτά Δεδομένα**: Δισεκατομμύρια τριπλέτες δεδομένων

Water, water, every where,
nor any drop to drink:
Samuel Taylor Coleridge,
Rime of the Ancient Mariner

Κλίμακα των Δεδομένων

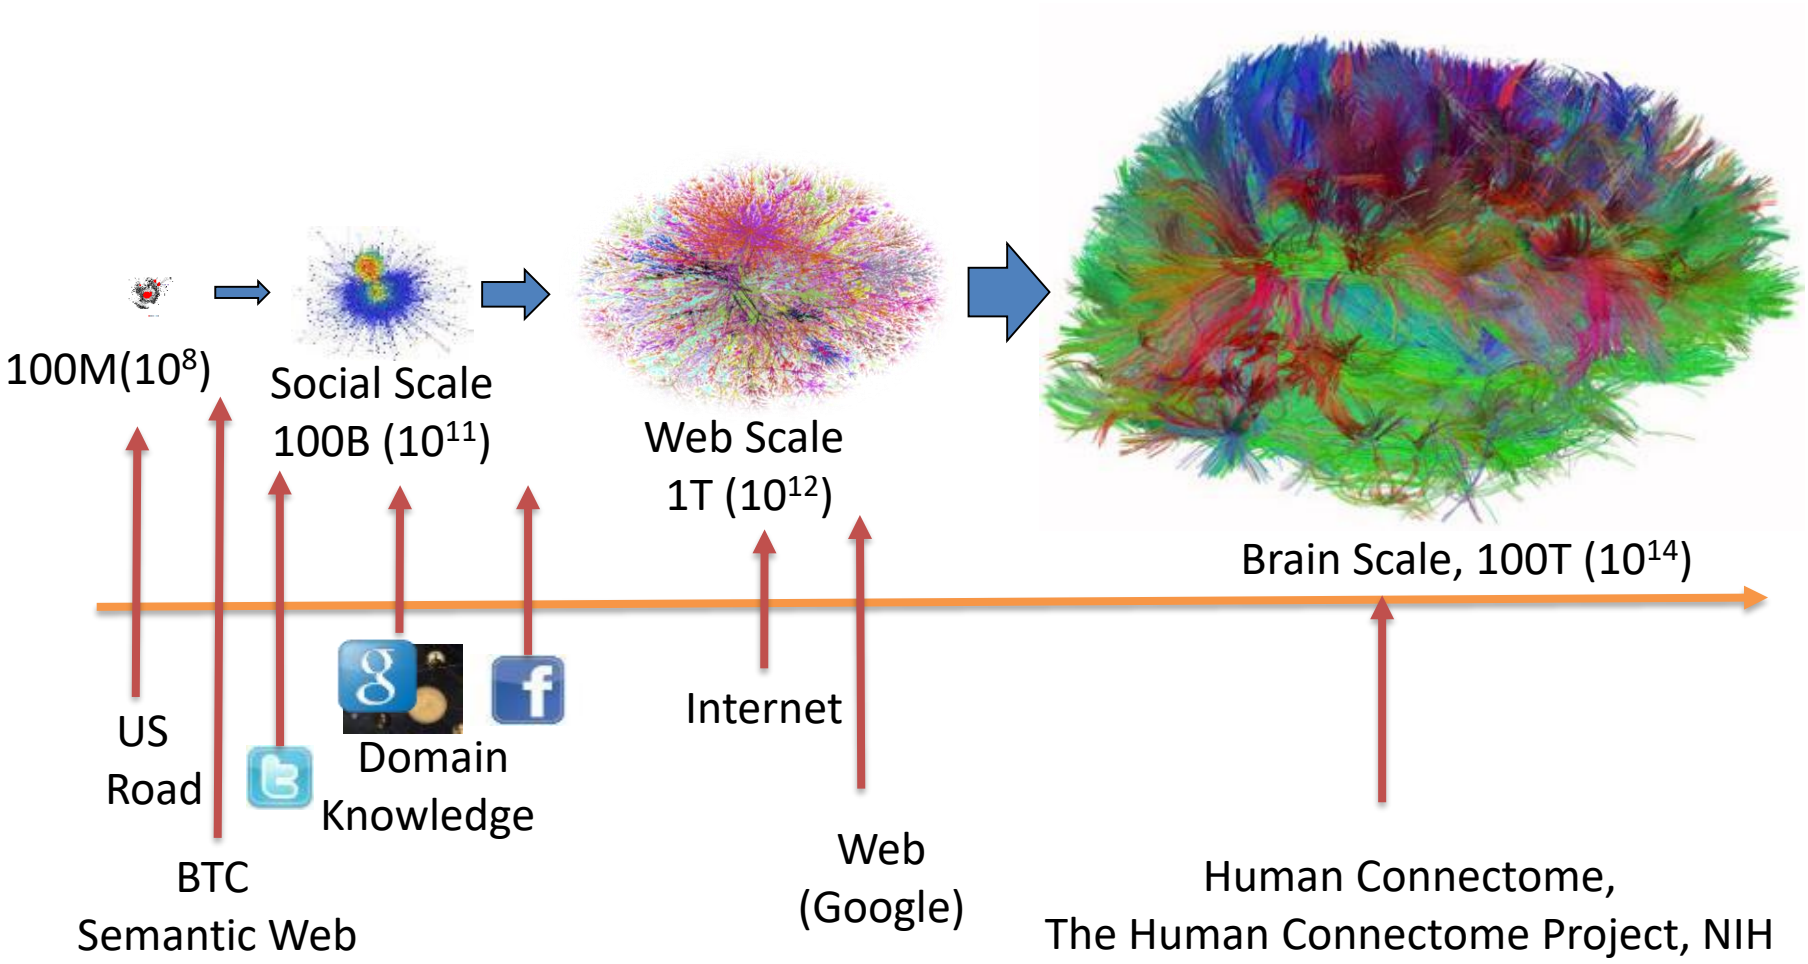

Acknowledgement: Y. Wu, WSU

Διασυνδεδεμένος Κόσμος

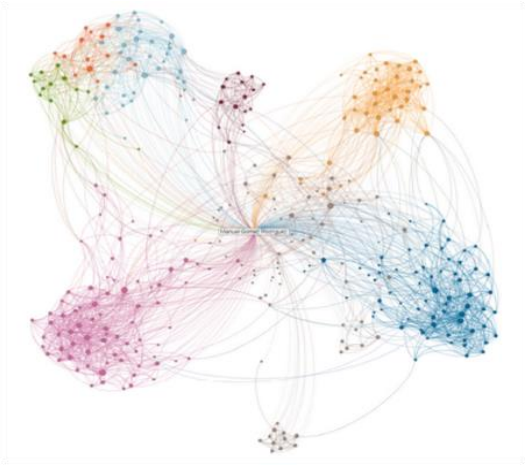

Social Networks

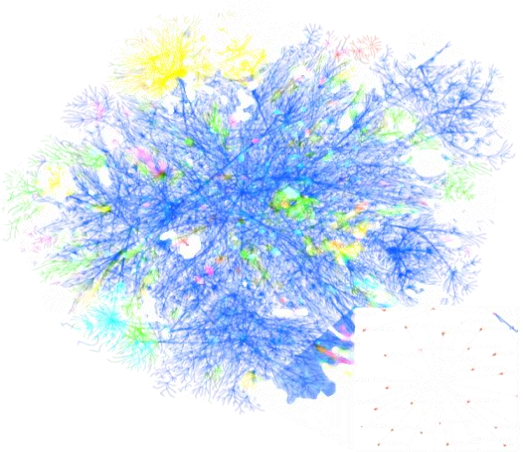

World Wide Web

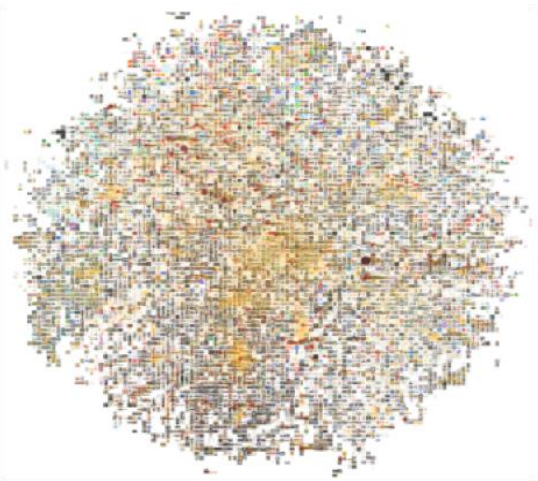

Information Networks

Transportation Networks

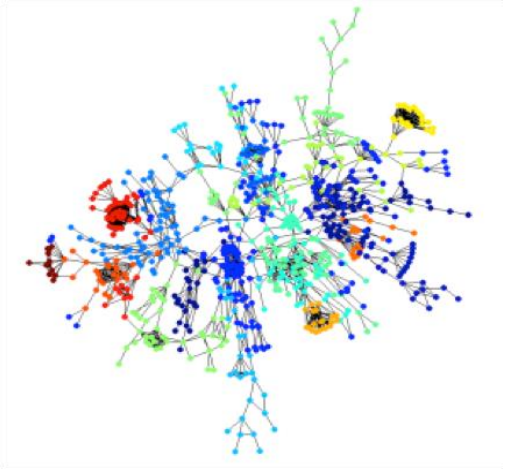

Protein Interactions

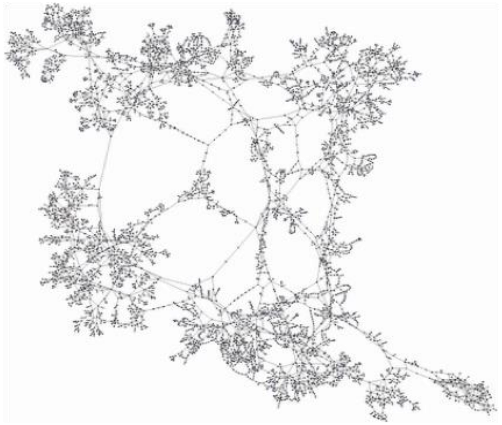

Electrical networks

Από κλασσικές Βάσεις Δεδομένων στα Μεγάλα Δεδομένα

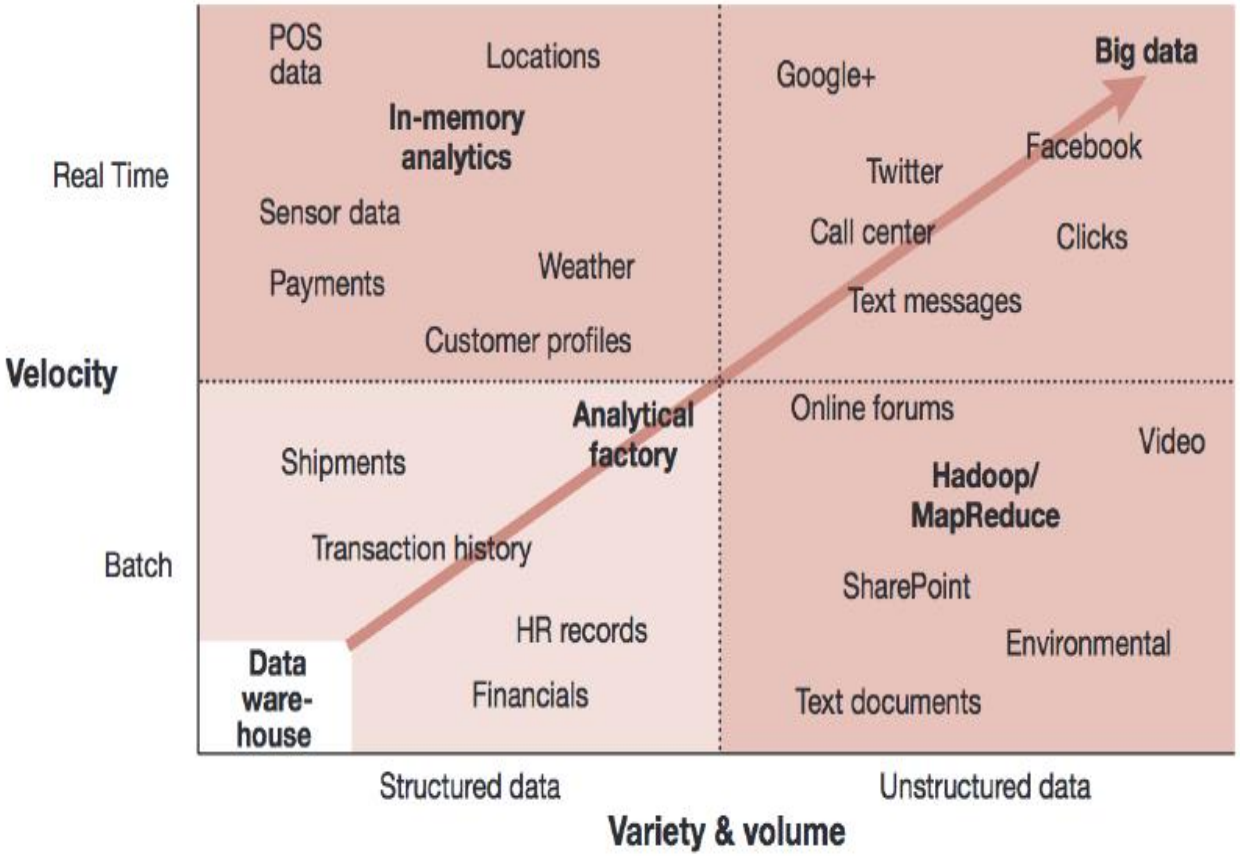

Source: Strategy& analysis

2. Τεχνητή Νοημοσύνη & Μηχανική Μάθηση

Μάθηση ↔ Νοημοσύνη

Νοημοσύνη είναι η δυνατότητα να μαθαίνουμε και να χρησιμοποιούμε εννοιολογικές προσεγγίσεις για να λύνουμε προβλήματα

Μηχανική Μάθηση ↔ Τεχνητή Νοημοσύνη

- *Η ΤΝ είναι η επιστήμη που επιτρέπει σε μηχανές να εκτελούν εργασίες που απαιτούν νοημοσύνη αν εκτελούνται από ανθρώπους* (Minsky 1986)
- *Η ΜΜ είναι η περιοχή της ΤΝ που αναπτύσσει τεχνολογίες που επιτρέπουν στις μηχανές να μαθαίνουν*

Γιατί Μηχανική Μάθηση? ↔ Γιατί Τεχνητή Νοημοσύνη?

- *Για να δημιουργούμε μηχανές με νοήμονα συμπεριφορά (δηλ. που είναι ικανές να συλλογίζονται, να προβλέπουν και να προσαρμόζονται) ενώ βοηθούν τους ανθρώπους στην εργασία τους, στην εκπαίδευσή τους, στην διασκέδασή τους.*

General Questions:

- Where are AI technologies used now & where will they come next?

- How is AI likely to change my job in the next ten years?

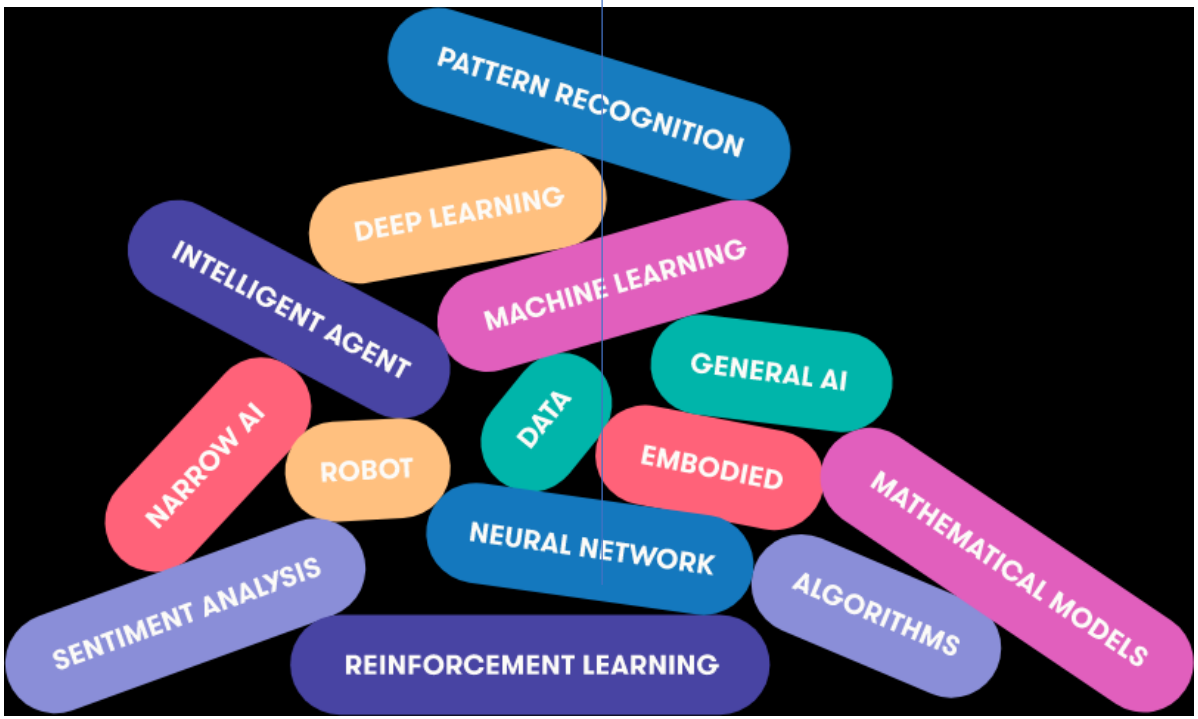

Question 1:

- But what is AI?

- It can be said that AI is bringing Natural Intelligence to machines.
- It is, however, more useful to list properties that are characteristic to AI systems:

Autonomy

The ability to perform tasks in complex environments without constant guidance by a user.

Adaptivity

The ability to improve performance by learning from experience.

Question 2:

- Which are the main AI subfields?

Machine learning

- Systems that improve their performance in a given task with more and more experience or data (Data driven)

Knowledge representation & Reasoning

- Machine understanding of rules, axioms, ontologies, semantic interoperability (Knowledge driven)

General vs narrow AI

- Narrow AI handles one task. General AI refers to a machine that can handle any intellectual task.

Question 3:

- Which are the main ML methodologies?

- **Supervised learning**: Being given an input, for example a photograph showing an animal, the system learns to predict the correct animal category → Humans provide many labeled examples of different animals.
- **Unsupervised learning:** Data are provided for training, but with no labels; the system groups similar inputs to form “clusters”, or reduces data to a small number of important “dimensions”.
- **Reinforcement learning:** An ML system, like a self-driving car, operates in an environment; feedback about good or bad choices is provided, by the environment, with some delay and used to update it.

Question 4:

- How do we train an ML system, in supervised learning, to predict the outcome in a problem?

Training & Test Datasets

- Split the data we possess in two parts: the **training** data and the **test** data.
- Train the ML system using only training data → create a model that predicts the output based on the input data
- Test if the system has learnt to solve the problem, i.e., if it can generalize to other data → apply it to the test data and evaluate its performance on these.
- If performance is not good → the problem can be: **underfitting** (need to train a more complex system) **overfitting** (stop learning because system learns to analyze the specific data & not solve the problem); use a third (**validation**) dataset to check when to stop & avoid overfitting.

Question 5:

- What is Deep Learning and Deep Neural Networks?

Deep Learning

- **DL** refers to ML techniques composed of many “layers” of simple processing units (neurons) connected in a way that the input to the system is passed through each layer in turn, producing the final output.
- **Deep Neural Networks** are such architectures, inspired by visual information processing in the brain.

Question 6:

- How do we train & use (Deep) Neural Networks?

1. Collection of Data (e.g., CIFAR-10)

Image Representation

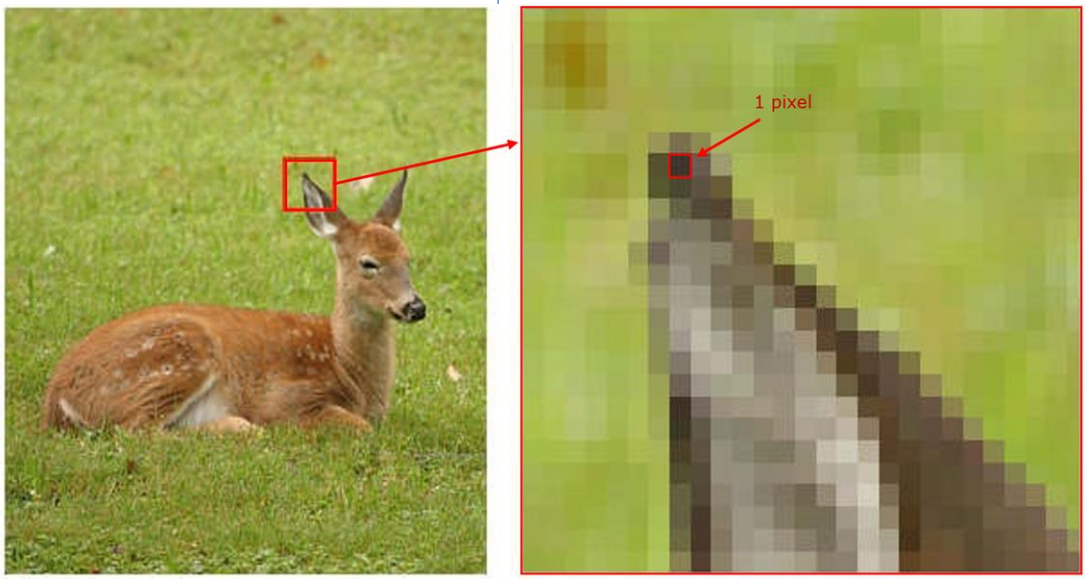

Question 6 (cont):

- How do we train (Deep) Neural Networks?

2. (Convolutional) Neural Networks, (C)NNs

- Input: the collected data/images (e.g., 3072 numerical values of the pixels per image)
- Output: the categories (e.g., 10 image classes)
- Number of Layers: e.g., 1 (hidden) layer of 100 neurons
- Training the (C)NN → compute the network parameters (weights: W1, W2) using the Backpropagation algorithm, the training dataset and the respective label.

Preview

[Zeiler and Fergus 2013]

Visualization of VGG-16 by Lane McIntosh. VGG-16 architecture from [Simonyan and Zisserman 2014].

Low-level features → Mid-level features → High-level features → Linearly separable classifier

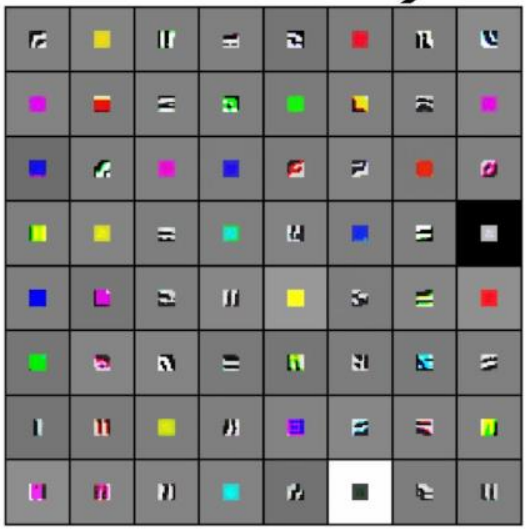

VGG-16 Conv1_1

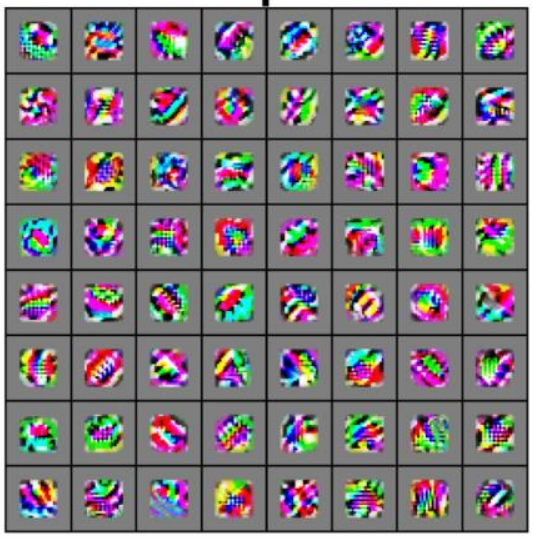

VGG-16 Conv3_2

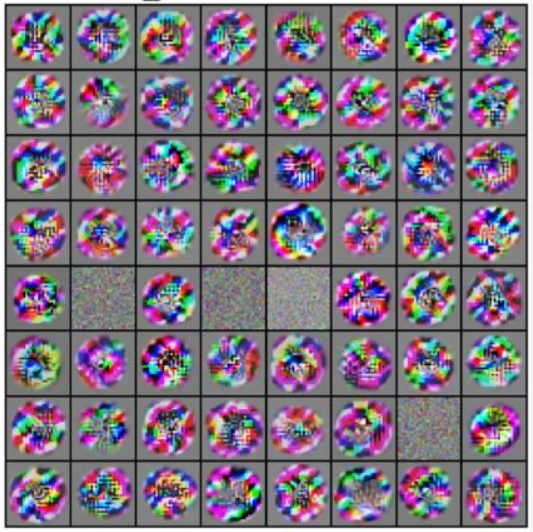

VGG-16 Conv5_3

Already known architectures (e.g. Neocognitron);
so, what changed? BIG DATA & Computing Power (GPUs)

- [Hubel & Wiesel 1962]:
  - simple cells detect local features
  - complex cells "pool" the outputs of simple cells within a retinotopic neighborhood.

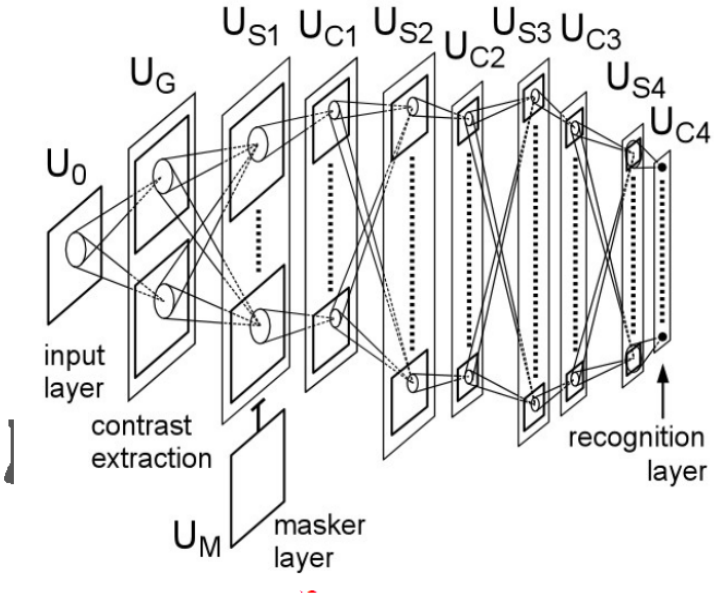

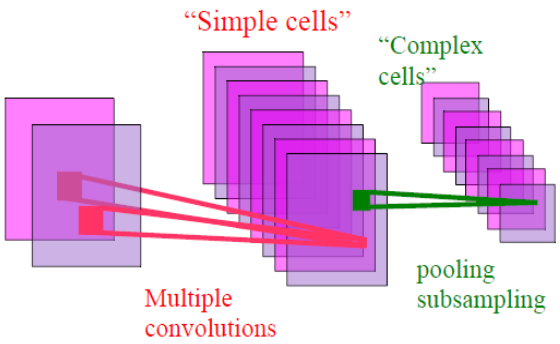

Question 7:

- How can we train a (C)NN for a specific application of ours?

- Do we need hundreds thousands of data?

TRANSFER LEARNING

- You do not need big amounts of data
- Many different NNs are available in the internet / github
- Use an existing trained NN for a similar/general type of problem and re-train part of it with your data.

Transfer Learning with CNNs

Donahue et al, "DeCAF: A Deep Convolutional Activation Feature for Generic Visual Recognition", ICML 2014
Razavian et al, "CNN Features Off-the-Shelf: An Astounding Baseline for Recognition", CVPR Workshops 2014

1. Train on Imagenet

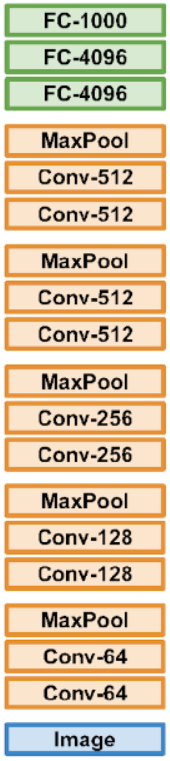

2. Small Dataset (C classes)

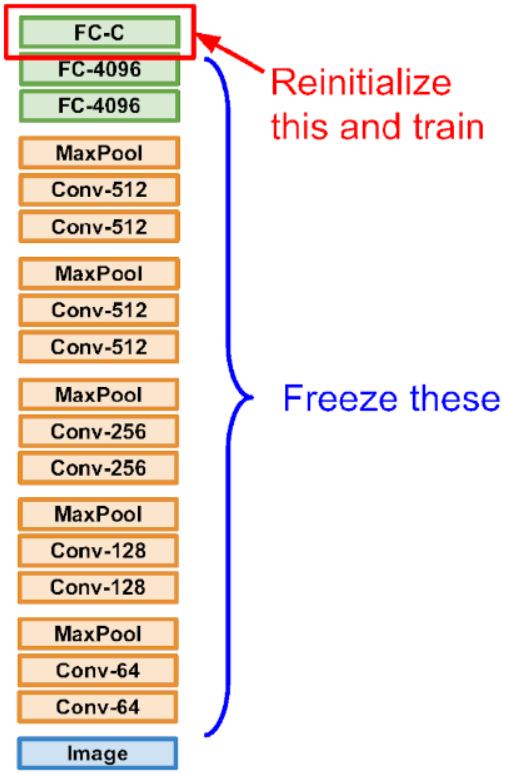

3. Bigger dataset

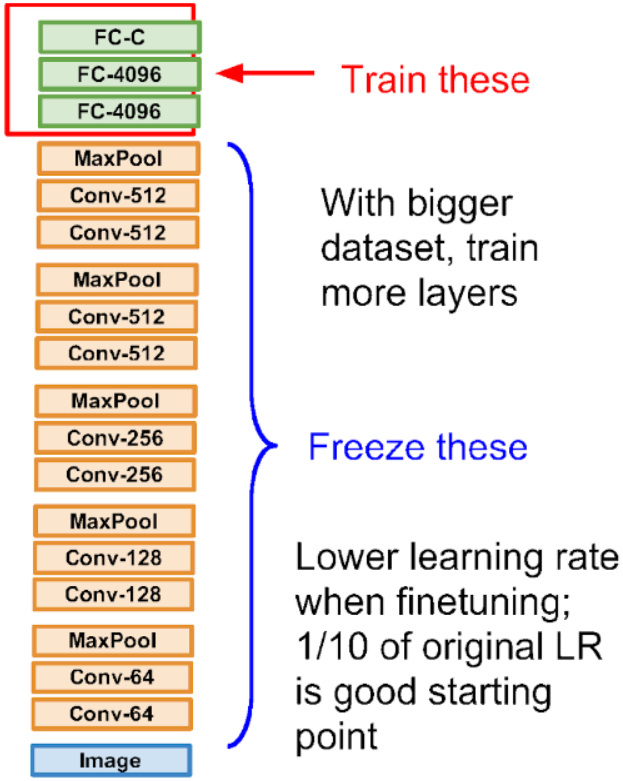

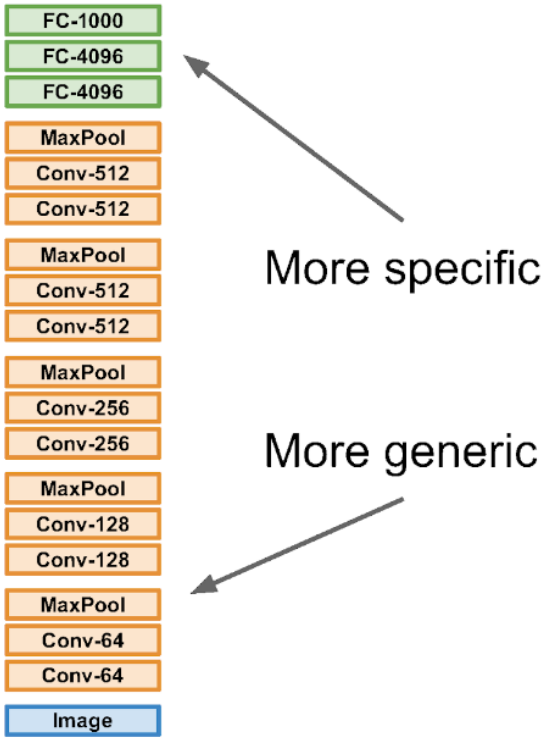

	very similar dataset	**very different dataset**
very little data	Use Linear Classifier on top layer	You're in trouble… Try linear classifier from different stages
quite a lot of data	Finetune a few layers	Finetune a larger number of layers

Question 8:

- What are GANs?

- What can we achieve with them?

GENERATIVE ADVERSARIAL NETWORKS

- Combine two CNNs, trained side by side & competing against each other
- One CNN is trained to generate images like the ones in a training data
- The other CNN is trained to separate the images generated by the first CNN from the real images in the training data
- Eventually the generated images are almost indistinguishable from real ones
- AND new, real-looking images are generated too!

GAN training

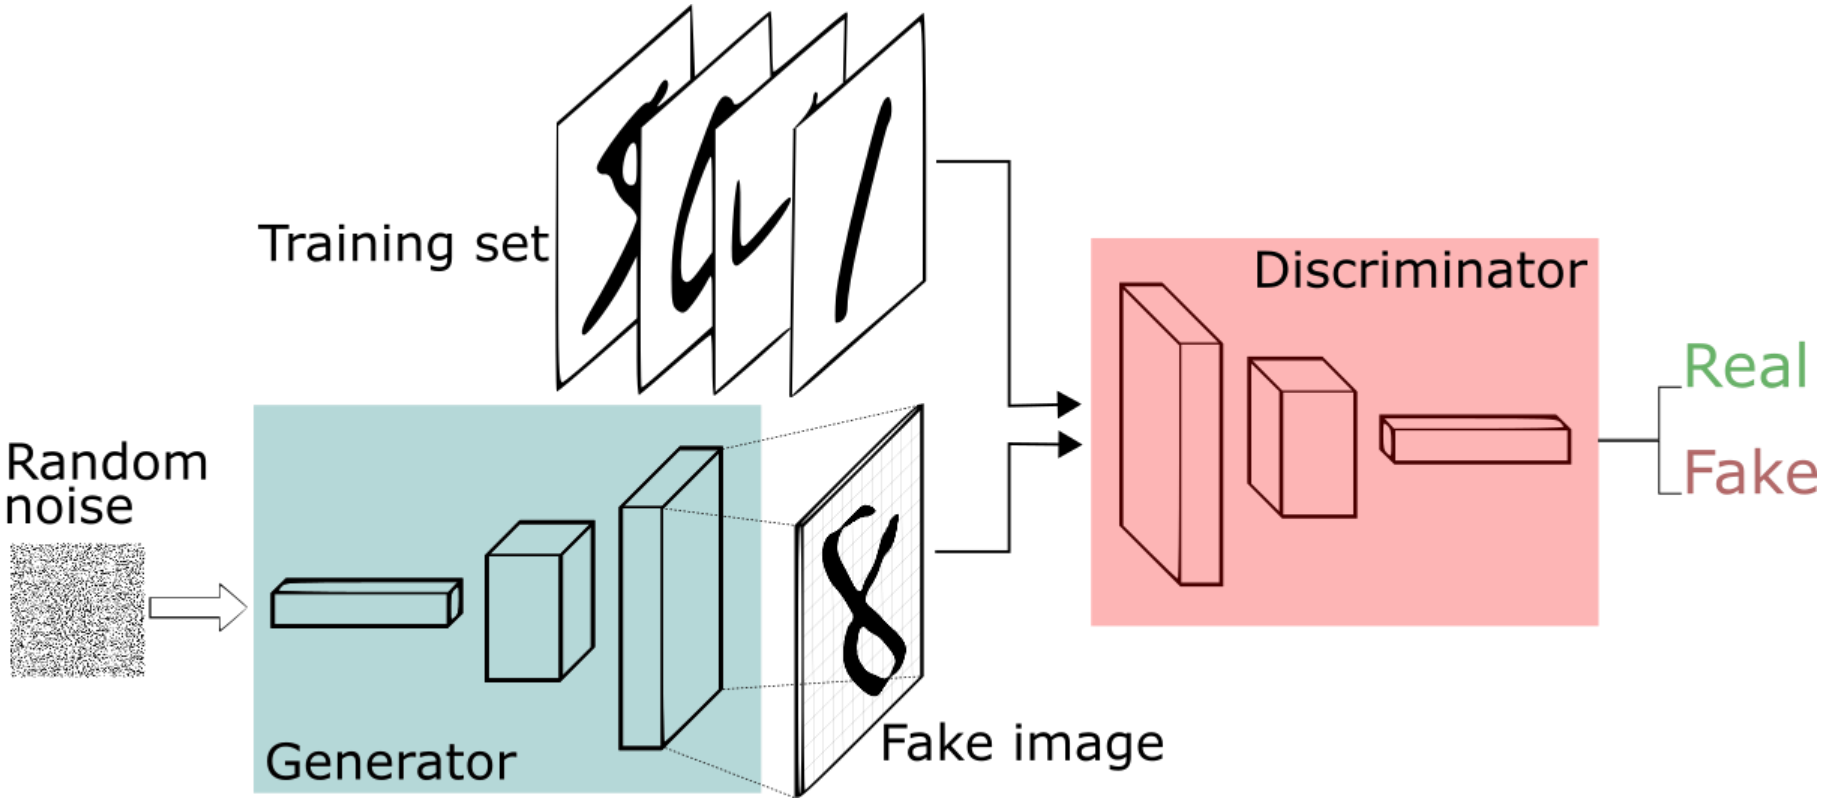

(Fake) images generated by GANs (by NVIDIA)

Final Question:

- How can we run ML experiments and use the outcomes?

Technical Requirements

- Digitization and organization of your data
- Definition of scenarios & targets (data analysis-supervised learning; new data generation-unsupervised learning; classification/regression/prediction?)
- Use of pretrained ML/DL models; retraining/ adaptation to your data and scenarios
- Use of a specific ML/DL infrastructure; computing resources: local server or cloud.
- Testing with different hyper-parameters, model dimensions, no of training iterations: auto-ML/AI techniques will assist you
- Evaluate results and continue experimenting.

The Strategic Importance of AI

Economic Potential of Artificial Intelligence, including Machine Learning

- Numerous Benefits: significant positive societal and economy-wide impacts, increase in productivity, innovation, growth and job creation (EU Parliament, 2020)
- AI is expected to underpin $15.7 trillion of global economic growth by 2030 (14% higher Global GDP as a result of AI) (PWC, 2017)

4. Ευφυής Αλληλεπίδραση Ανθρώπου Μηχανής

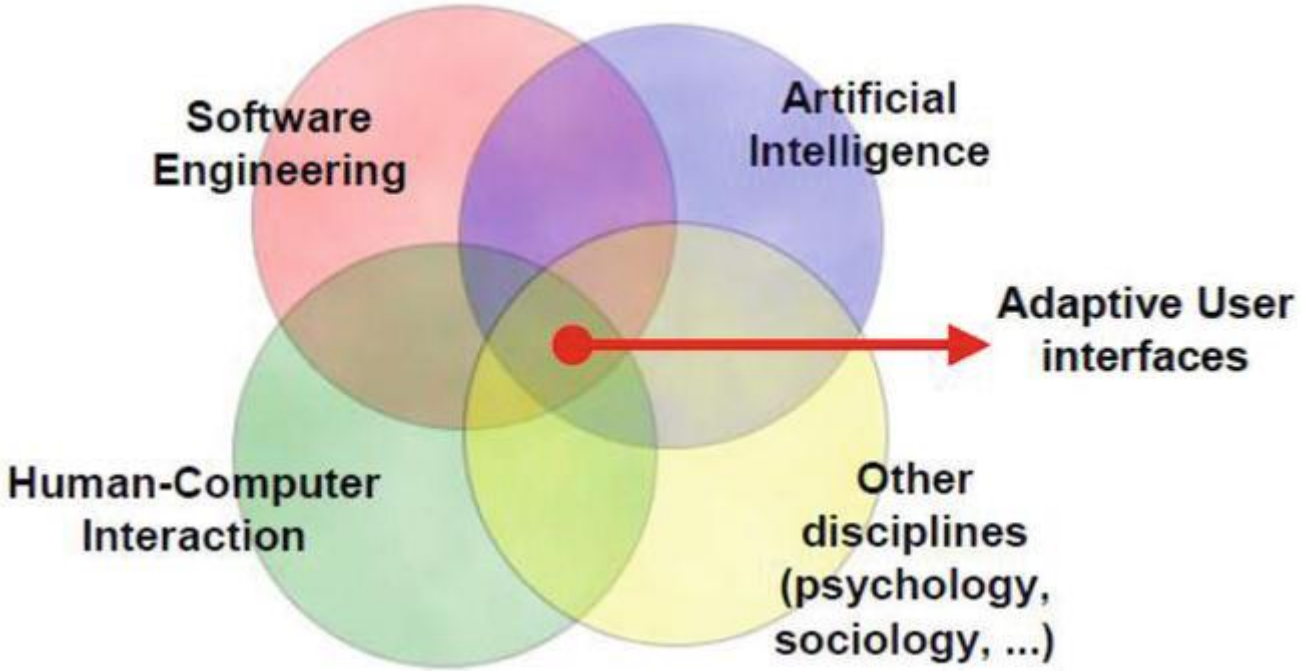

Ανθρωποκεντρική Σχεδίαση

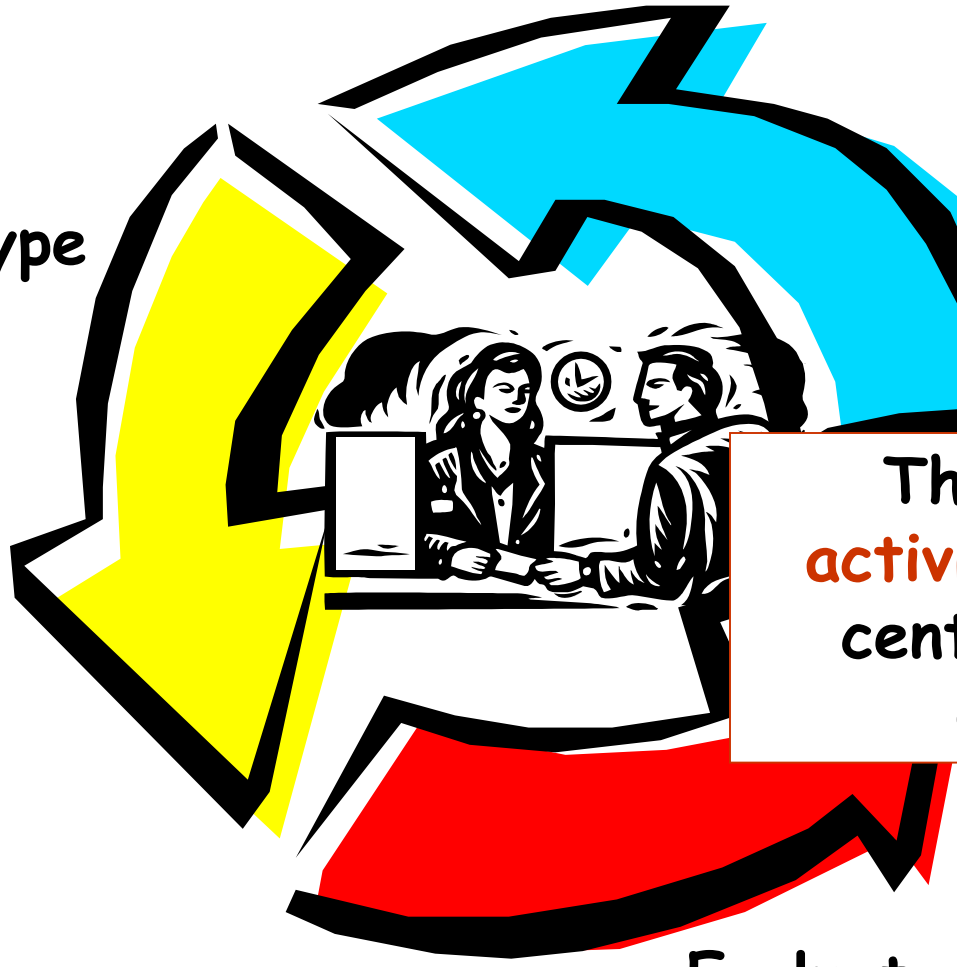

The **human activities** in the center of our design

Αναλυτική Μεγάλων Δεδομένων

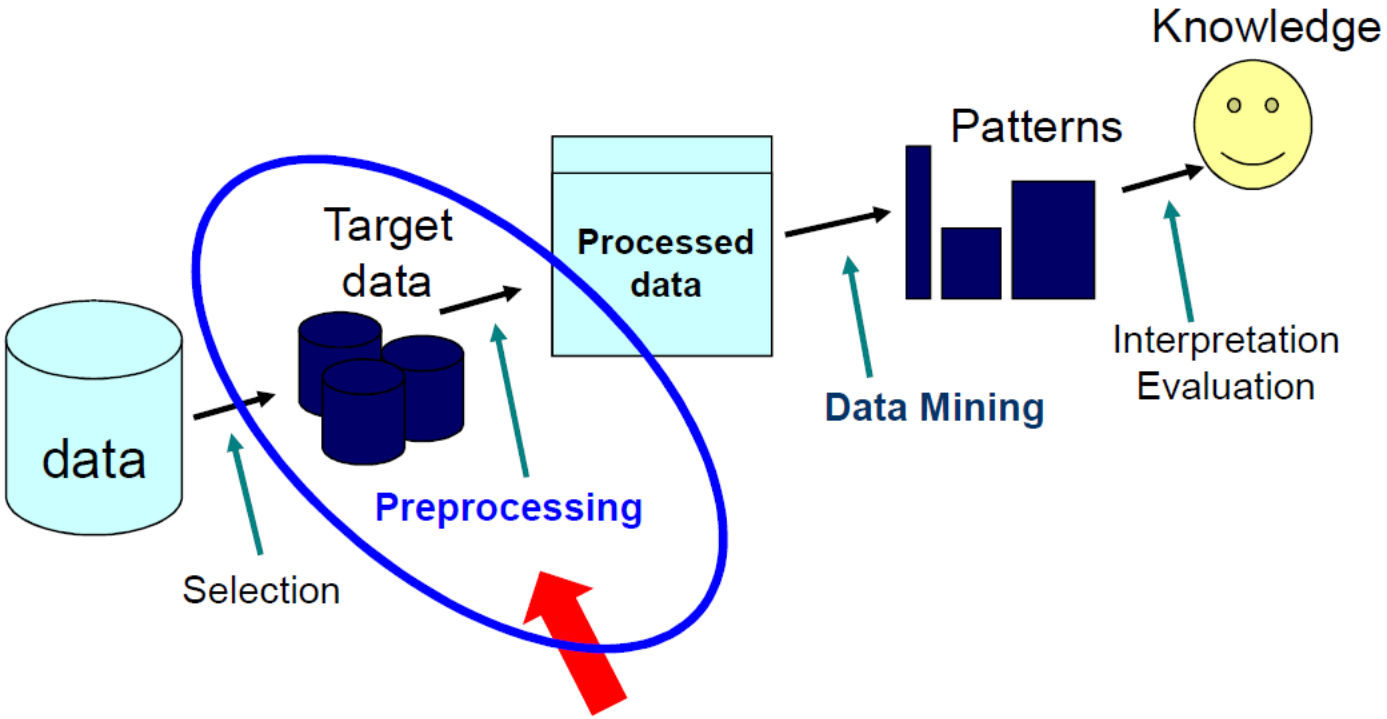

Στόχοι της Αναλυτικής Δεδομένων

Industry 4.0 to Society 5.0

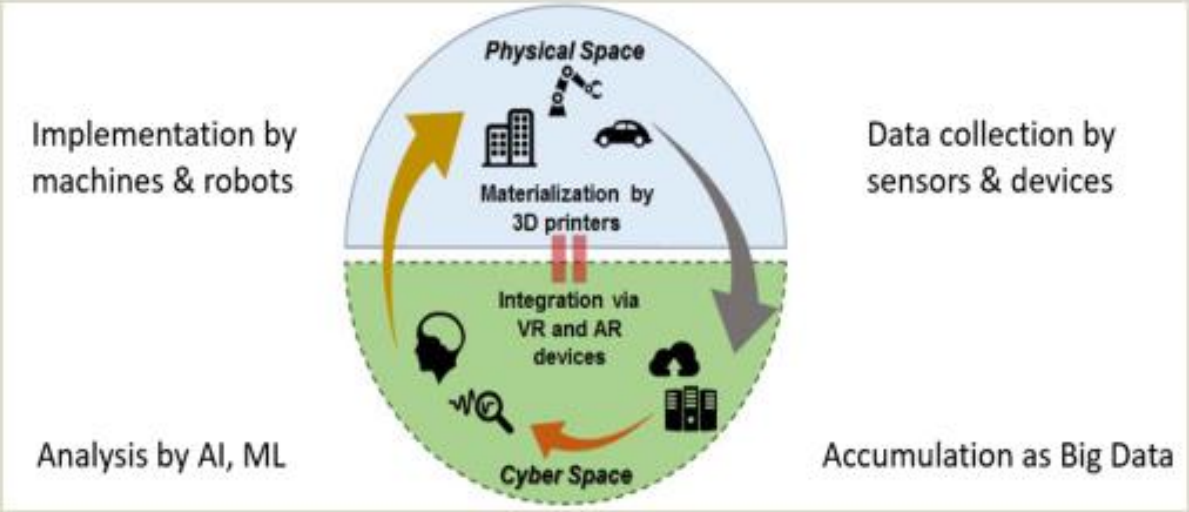

Ευφυής ΑΑΥ

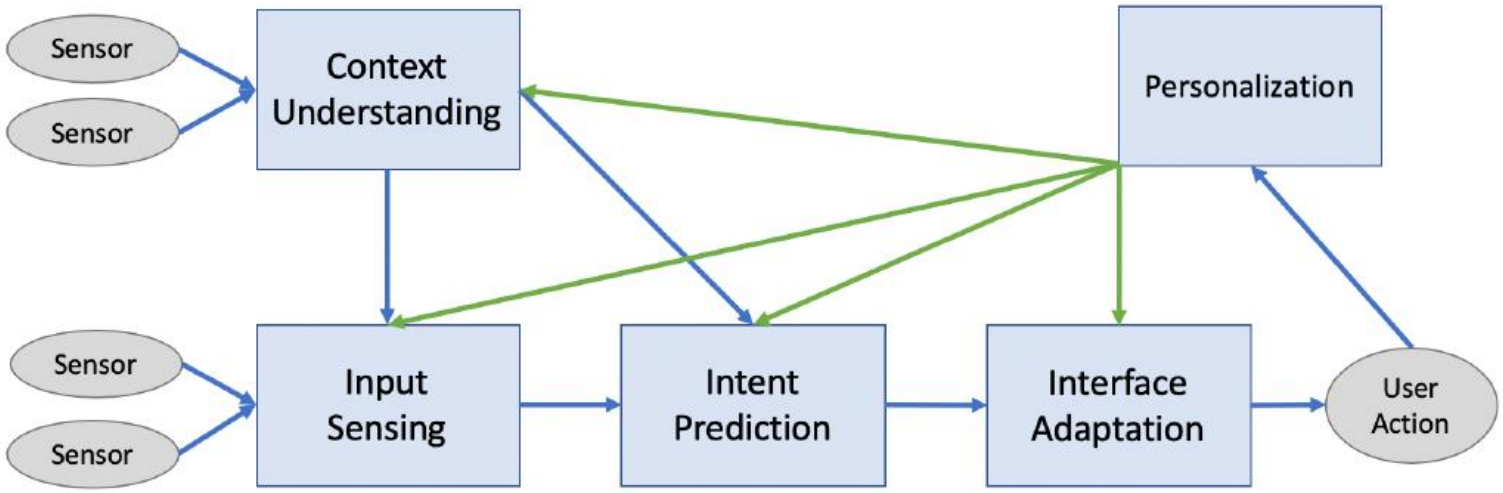
Sensor
Sensor
Context Understanding
Personalization
Sensor
Sensor
Input Sensing
Intent Prediction
Interface Adaptation
User Action

Μηχανική Μάθηση και ΑΑΥ

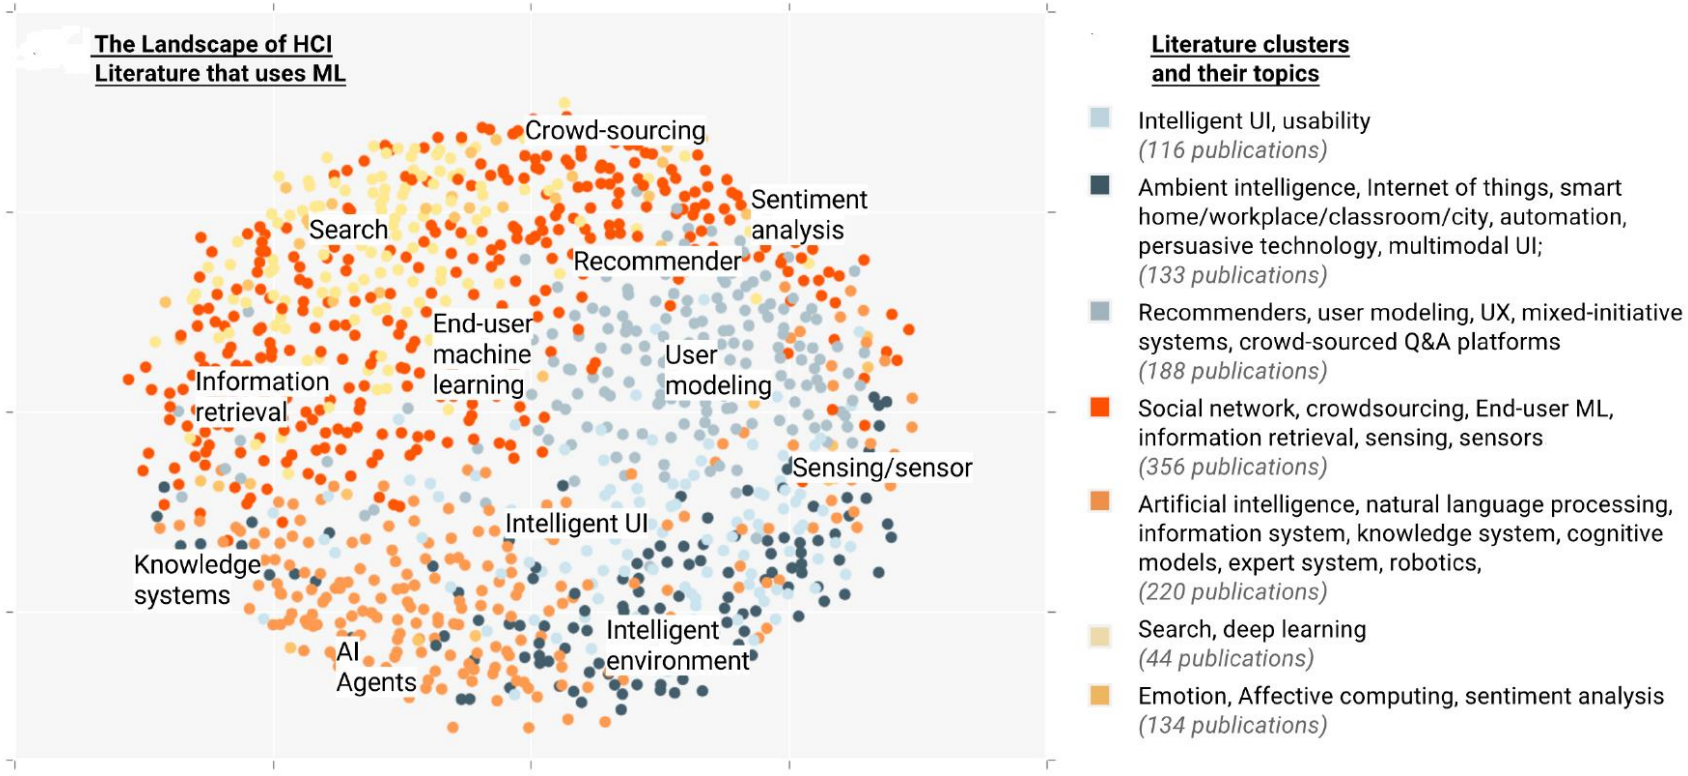

Μηχανική Μάθηση ←→ ΑΑΥ

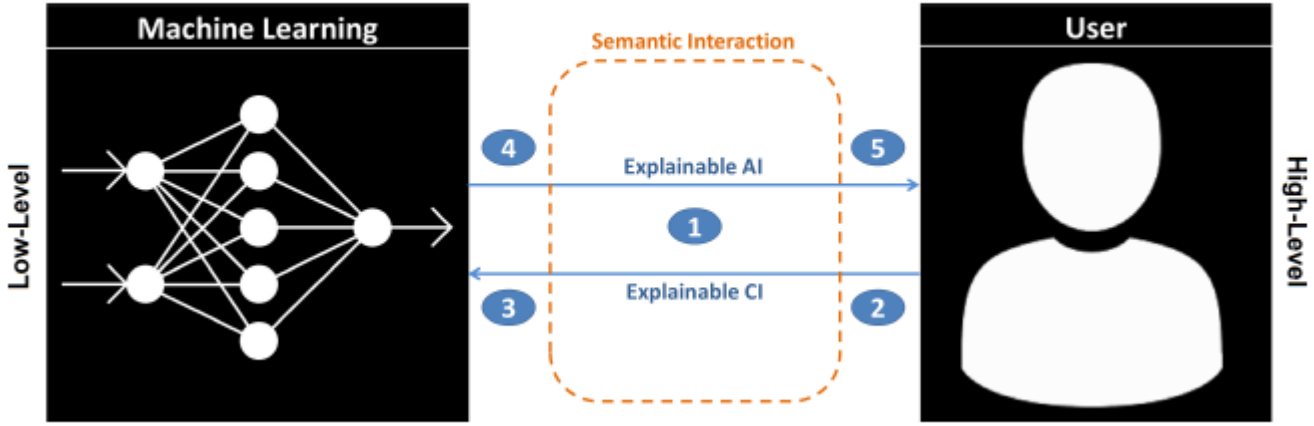

The human and AI black boxes are connected via two communication channels by Semantic Interaction (SI)

- Explainable Cognitive Intelligence (XCI) provides information about the human to the AI (2–3)
- Explainable Artificial Intelligence (XAI) provides information about the AI to the human (4-5)