# Machine Learning and Knowledge Extraction

## Introduction to CD-MAKE

**Andreas Holzinger**

Holzinger-Group, HCI-KDD, Institute for Medical Informatics, Statistics and Documentation, Medical University Graz
&
Institute of Interactive Systems and Data Science, Graz University of Technology

a.holzinger@hci-kdd.org
http://hci-kdd.org

- 01 What is the HCI-KDD approach?
- 02 Application Area: Health
- 03 automatic ML (aML)
- 04 interactive ML (iML)

# 01 What is the HCI-KDD approach?

- **ML is a very practical field – algorithm development is at the core – however, successful ML needs a concerted effort of various topics …**

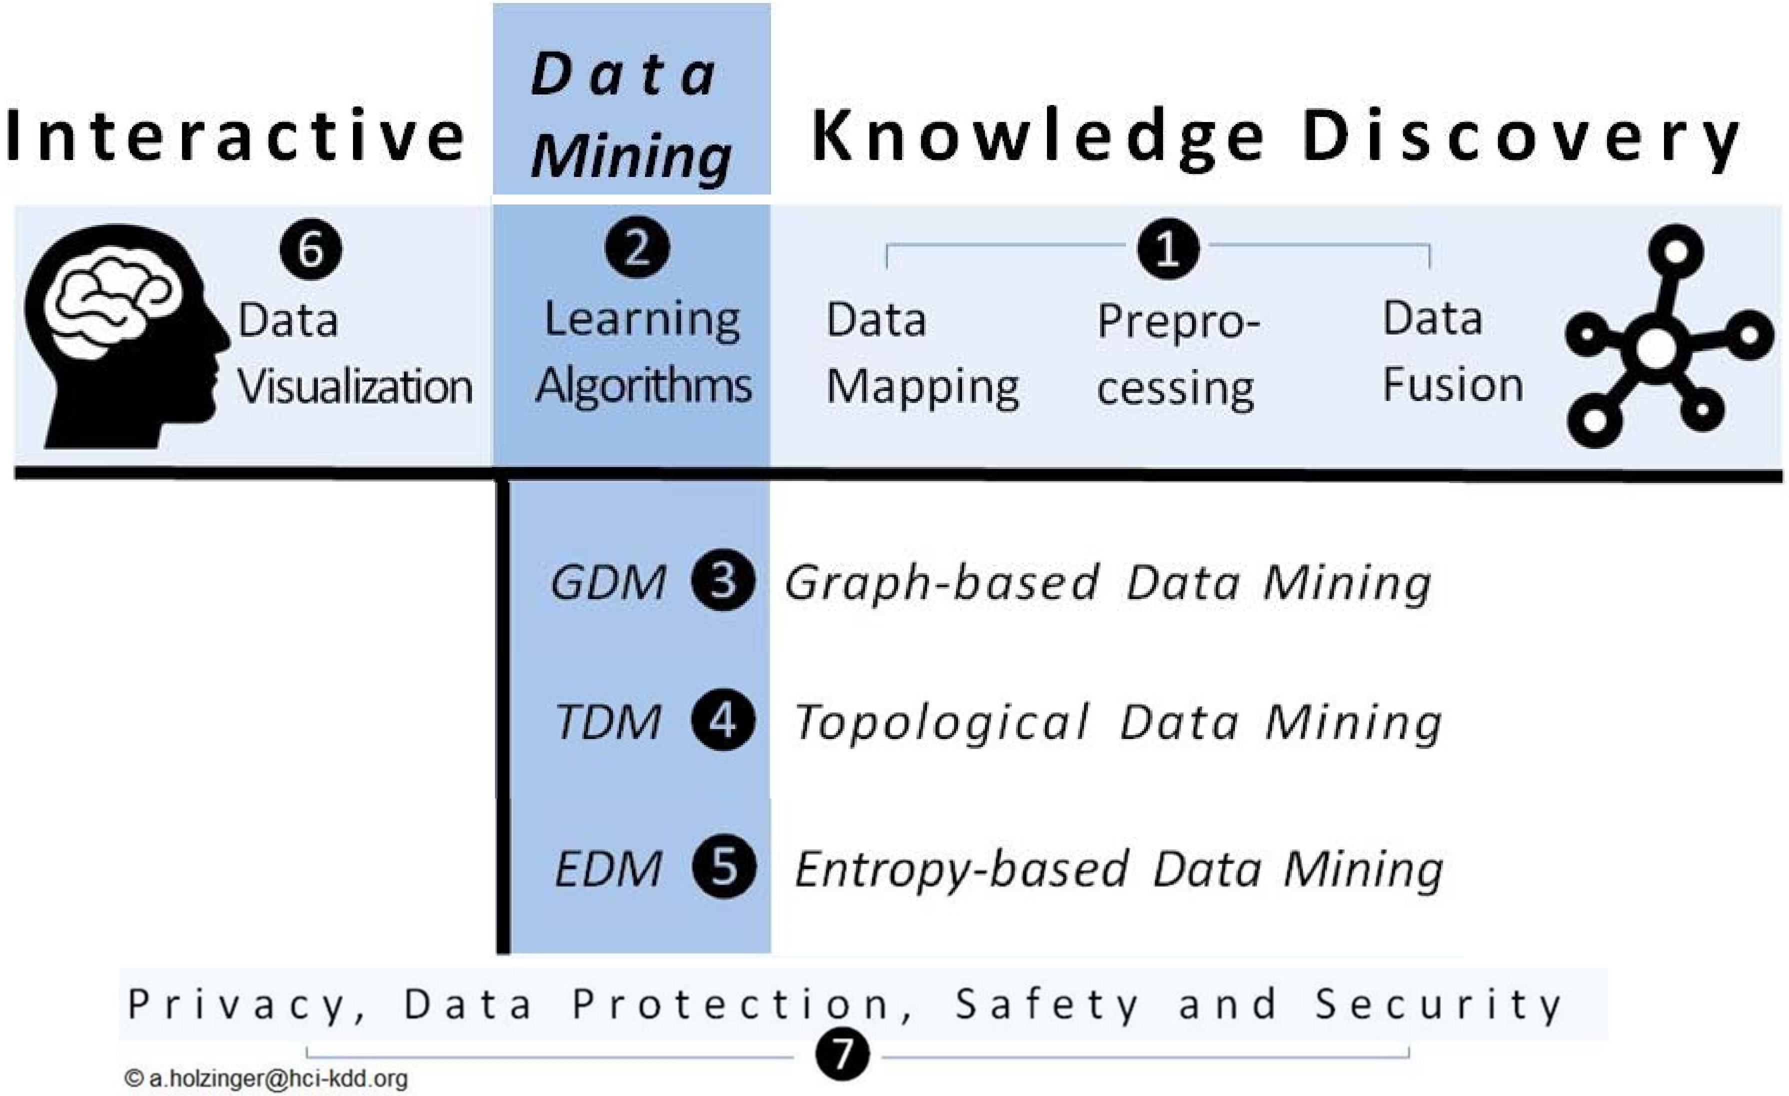


Holzinger, A. 2014. Trends in Interactive Knowledge Discovery for Personalized Medicine: Cognitive Science meets Machine Learning. IEEE Intelligent Informatics Bulletin, 15, (1), 6-14.


http://www.bach-cantatas.com


http://hci-kdd.org/international-expert-network


Holzinger, A. 2017. Introduction to Machine Learning and Knowledge Extraction (MAKE). Machine Learning and Knowledge Extraction, 1, (1), 1-20, doi:10.3390/make1010001

# Solve Intelligence then solve everything else

# "Solve intelligence – then solve everything else"


https://youtu.be/XAbLn66iHcQ?t=1h28m54s

**Demis Hassabis, 22 May 2015**

The Royal Society,
Future Directions of Machine Learning Part 2

Google DeepMind


- **Cognitive Science $\rightarrow$ human intelligence**
- **Computer Science $\rightarrow$ computational intelligence**
- **Human-Computer Interaction $\rightarrow$ the bridge**

Holzinger, A. 2013. Human–Computer Interaction and Knowledge Discovery (HCI-KDD): What is the benefit of bringing those two fields to work together? In: LNCS 8127. Springer, pp. 319-328, doi:10.1007/978-3-642-40511-2_22.

- 1) **learn** from prior data
- 2) **extract** knowledge
- 2) **generalize**, i.e. guessing where a probability mass function concentrates
- 4) fight the curse of **dimensionality**
- 5) **disentangle** underlying explanatory factors of data, i.e.
- 6) **understand** the data in the **context** of an application domain

# 02 Application Area Health Informatics

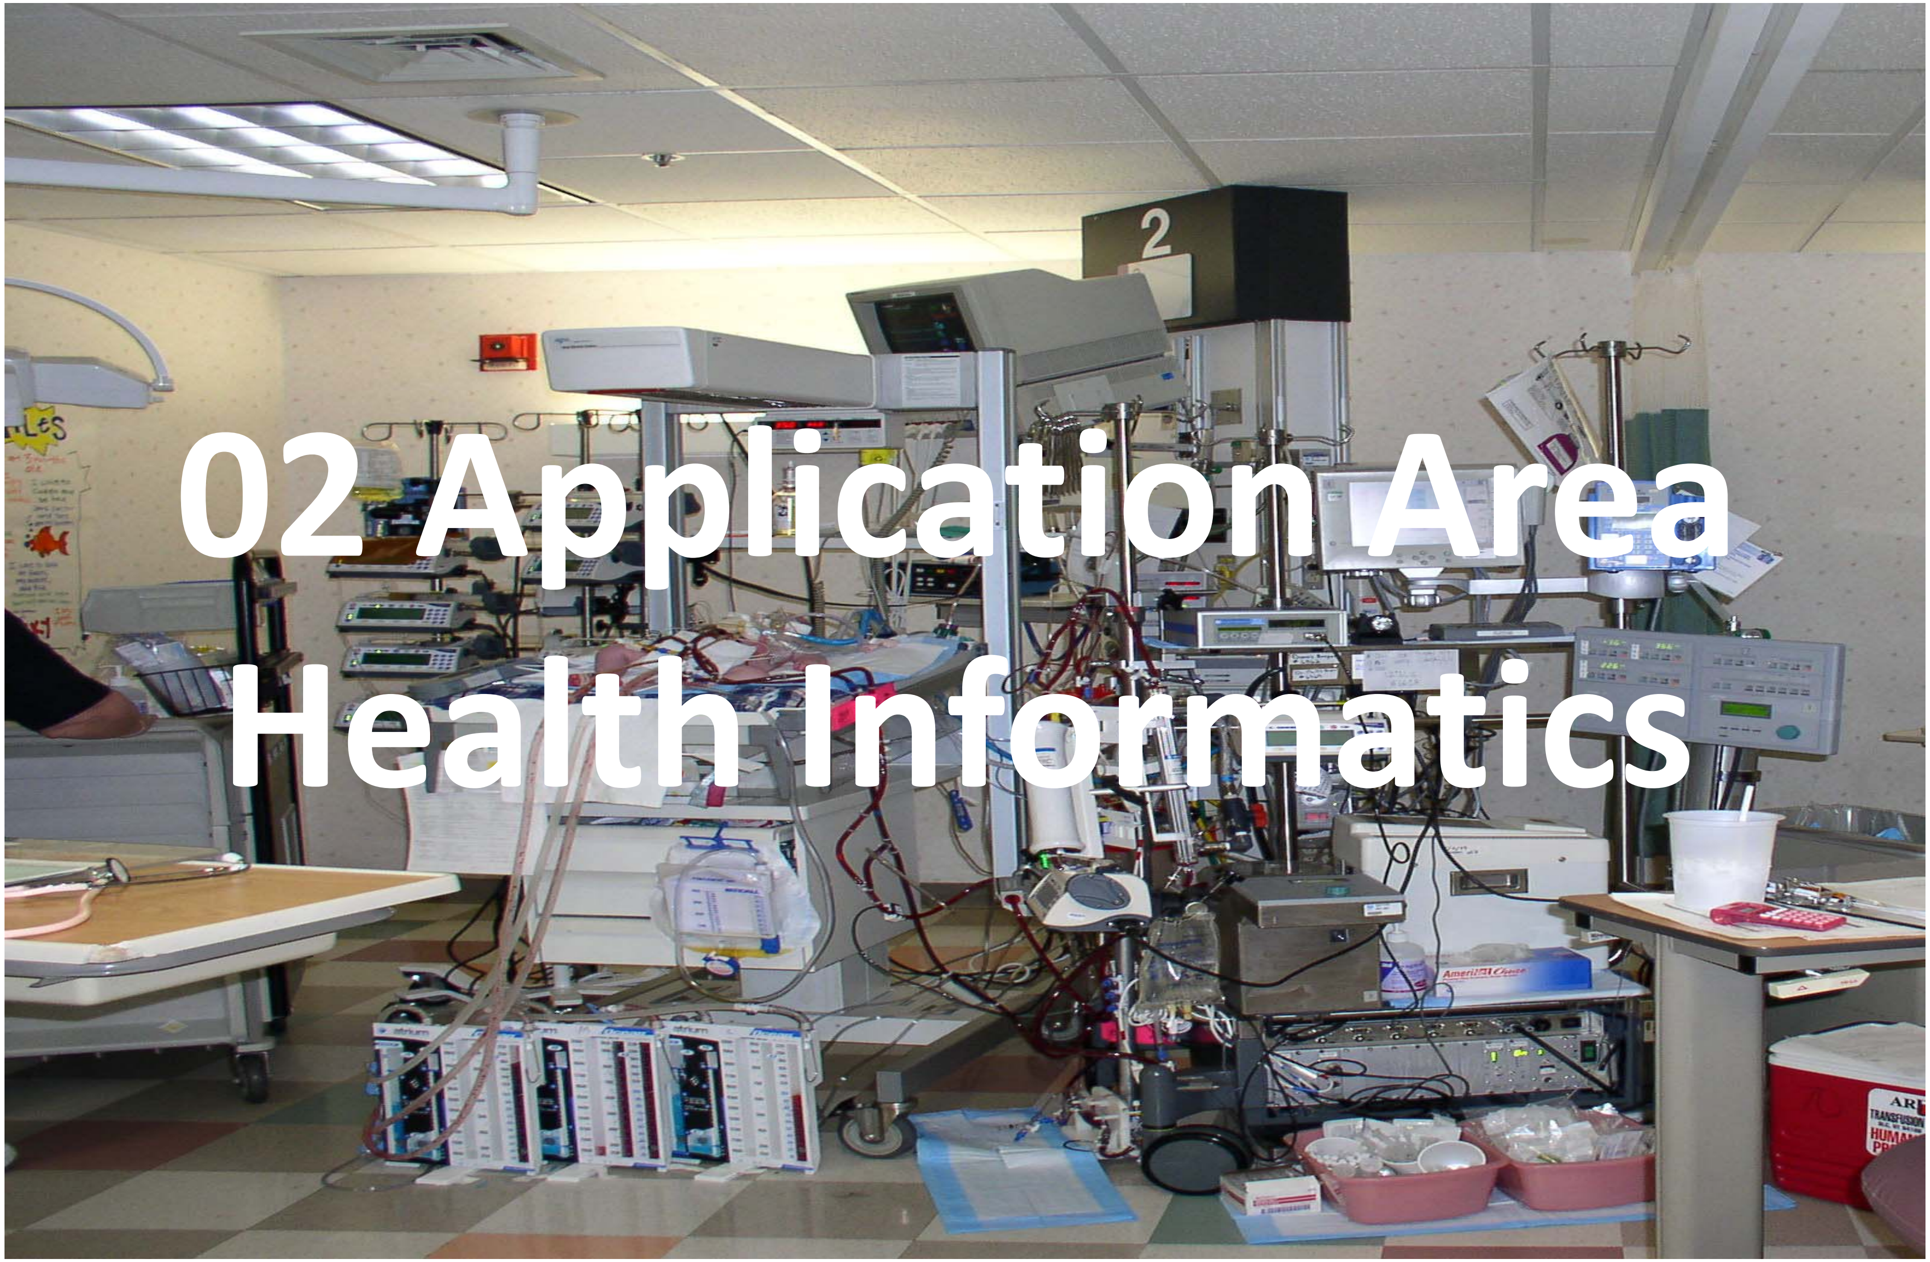

# Why is this application area complex ?

# In medicine we have two different worlds …



# Our central hypothesis:
# Information may bridge this gap

Holzinger, A. & Simonic, K.-M. (eds.) 2011. *Information Quality in e-Health. Lecture Notes in Computer Science LNCS 7058, Heidelberg, Berlin, New York: Springer.*


Where is the
problem in building
this bridge?

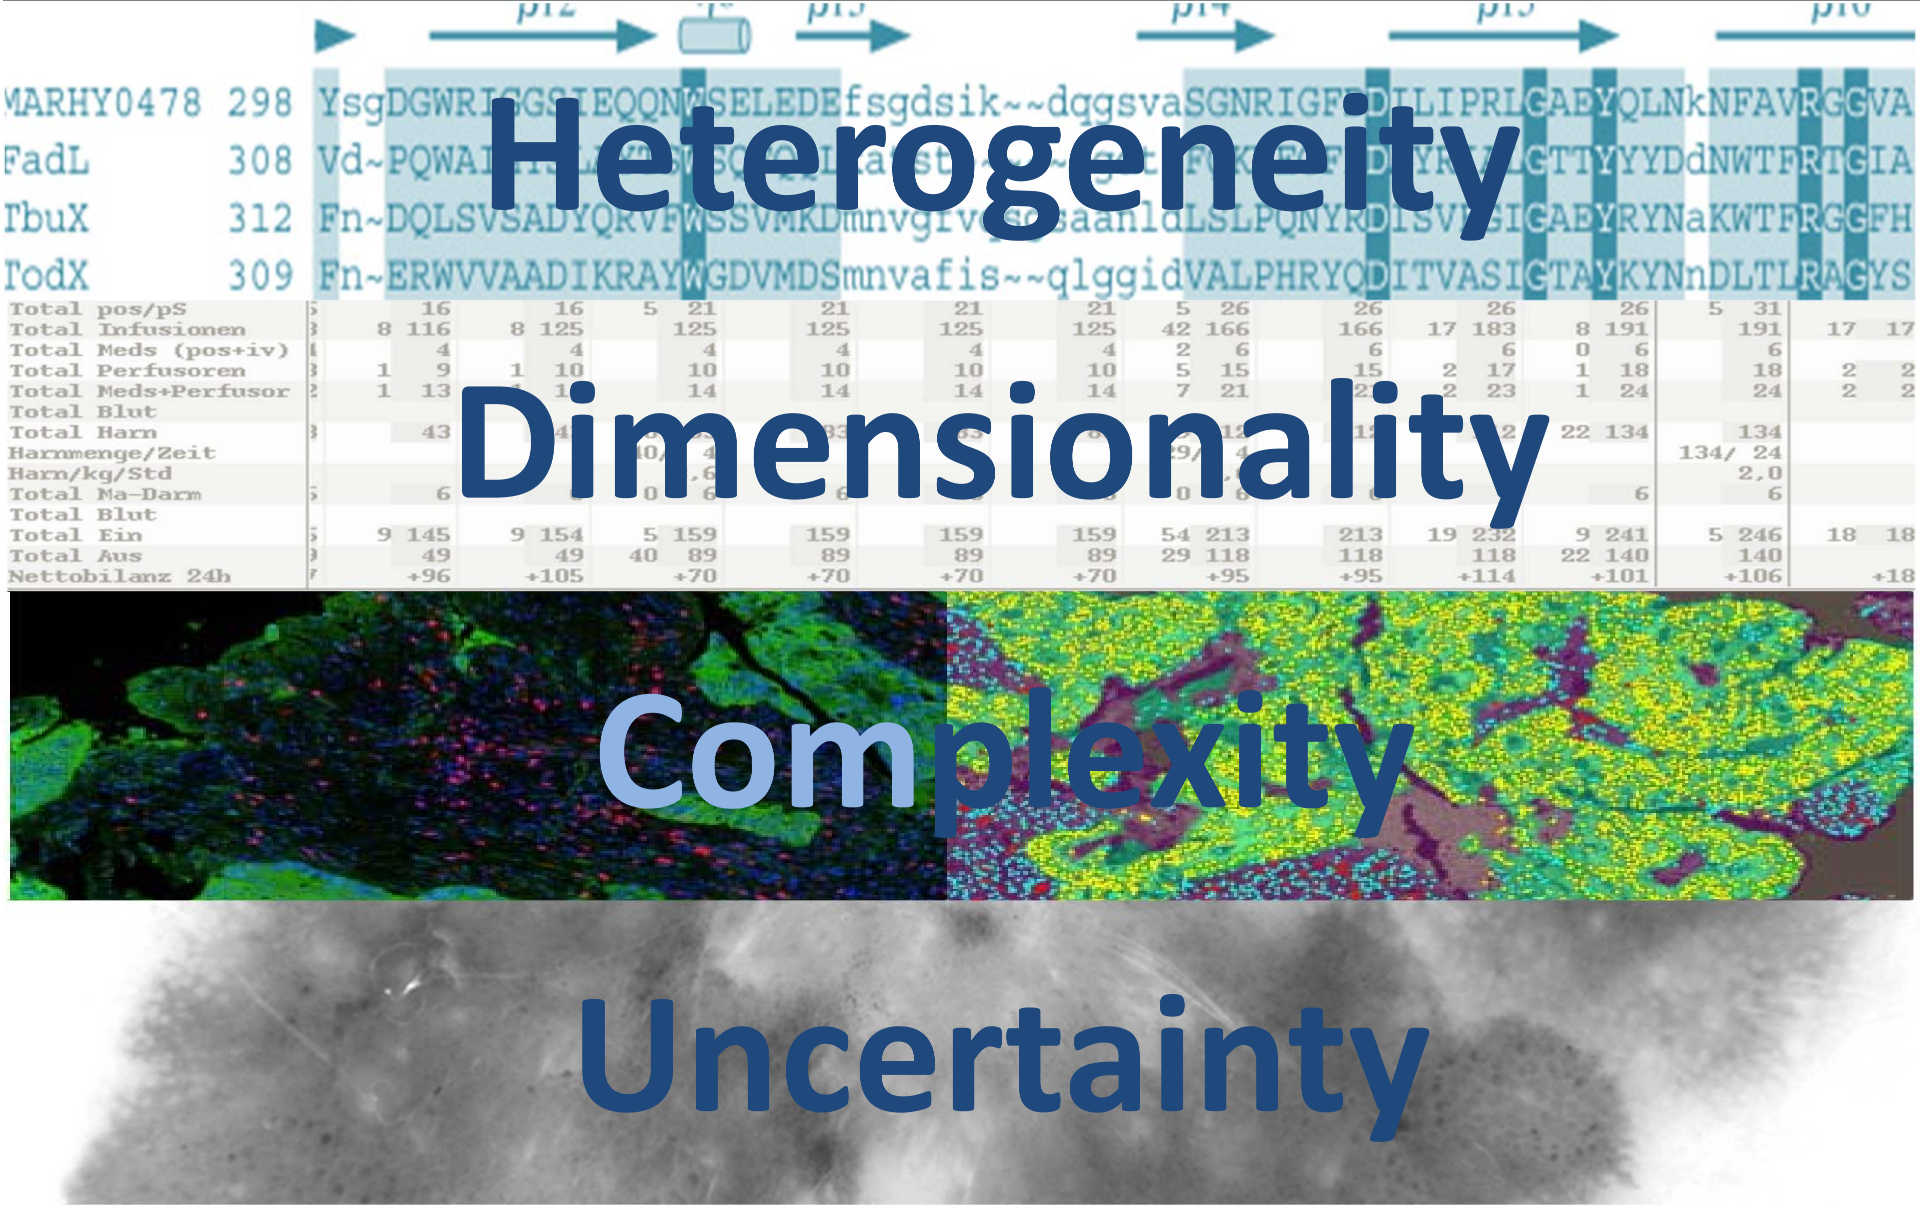


Holzinger, A., Dehmer, M. & Jurisica, I. 2014. Knowledge Discovery and interactive Data Mining in Bioinformatics - State-of-the-Art, future challenges and research directions. BMC Bioinformatics, 15, (S6), I1.

# 03 Probabilistic Information p(x)

# Learning and Inference

$d$ ... *data*

$h$ ... *hypotheses*

$\mathcal{H}$ ... $\{H_1, H_2, \dots, H_n\}$ $\quad \forall\, h, d$ ...

Likelihood

Prior Probability

$$p(h|d) = \frac{p(d|h) * p(h)}{\sum_{h \in H} p(d|h')\, p(h')}$$

Posterior Probability

Problem in $\mathbb{R}^n \rightarrow$ complex

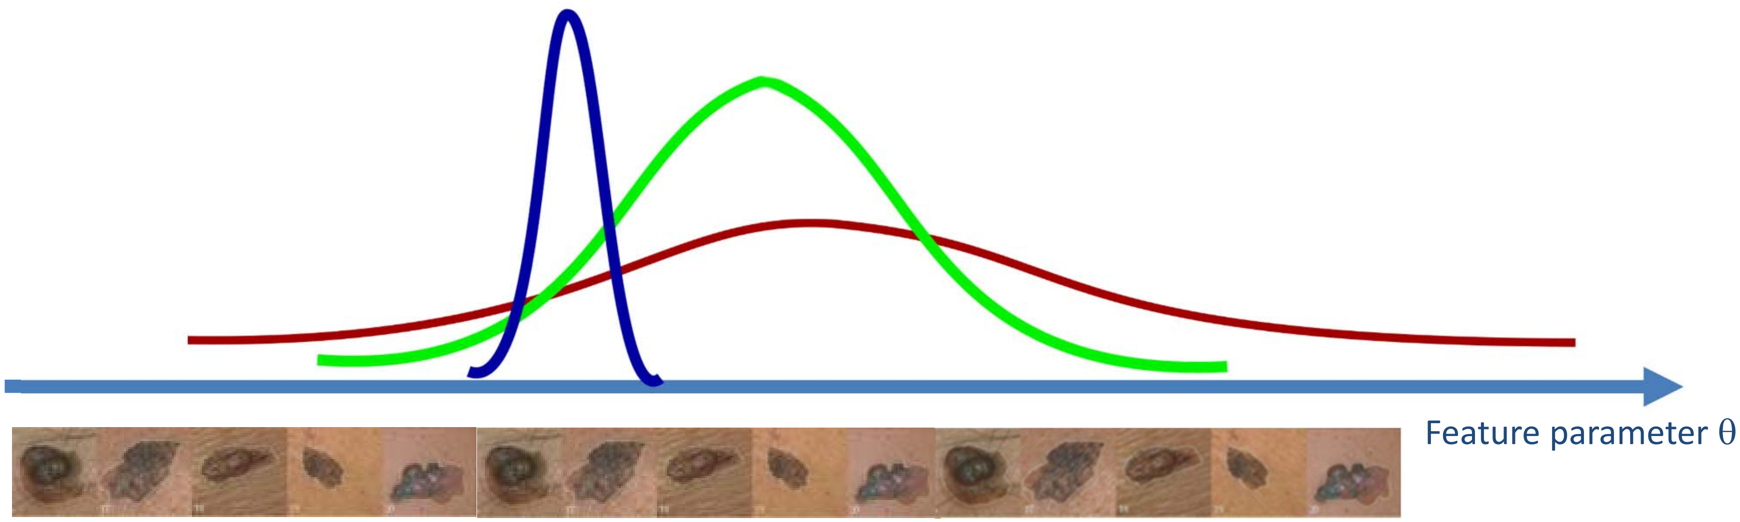

$$\mathcal{D} = x_{1:n} = \{x_1, x_2, ..., x_n\}$$

$$p(\mathcal{D}|\theta)$$

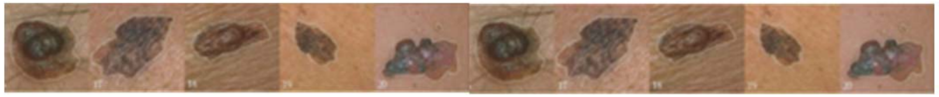

$$p(\theta|\mathcal{D}) = \frac{p(\mathcal{D}|\theta) * p(\theta)}{p(\mathcal{D})}$$

**The inverse probability allows to learn from data, infer unknowns, and make predictions**

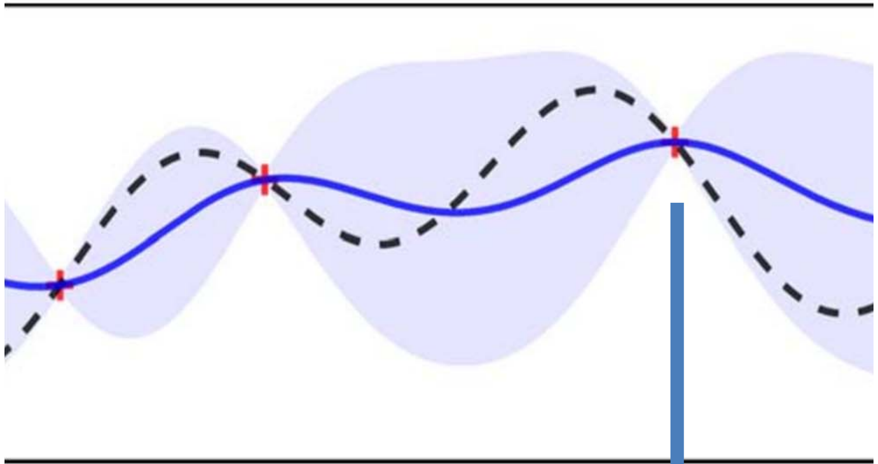

**Algorithm 1** Bayesian optimization

**for** $n = 1, 2, \ldots$ **do**
- select new $\mathbf{x}_{n+1}$ by optimizing acquisition function $\alpha$

$$\mathbf{x}_{n+1} = \arg\max_{\mathbf{x}} \alpha(\mathbf{x}; \mathcal{D}_n)$$

- query objective function to obtain $y_{n+1}$
- augment data $\mathcal{D}_{n+1} = \{\mathcal{D}_n, (\mathbf{x}_{n+1}, y_{n+1})\}$
- update statistical model

**end for**

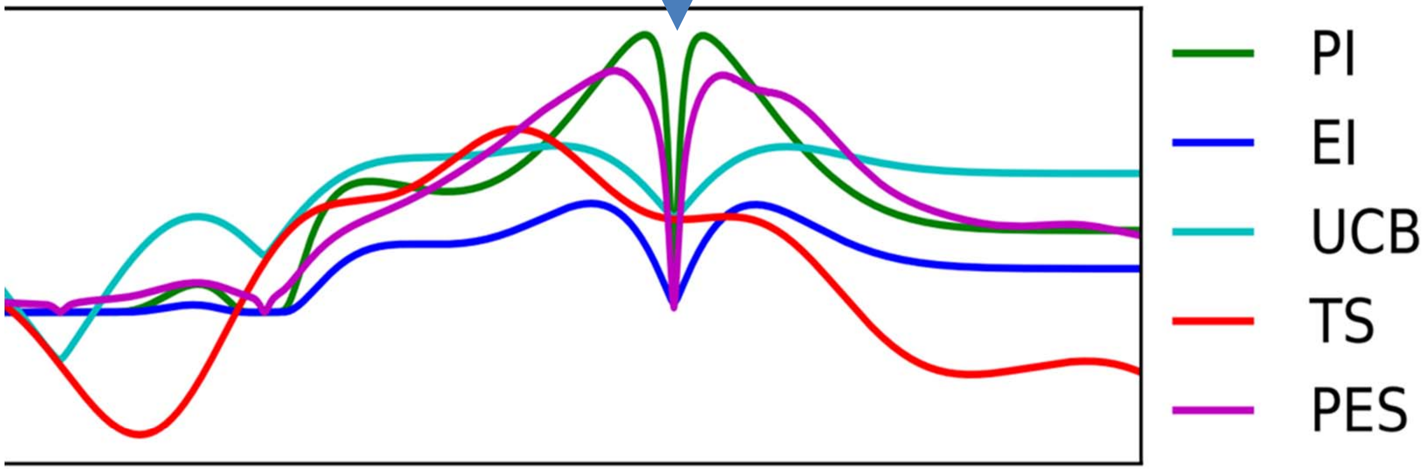


Probability of Improvement

Expected Improvement

Upper Confidence Bound

Thompson Sampling

Predictive Entropy Search

Shahriari, B., Swersky, K., Wang, Z., Adams, R. P. & De Freitas, N. 2016. **Taking the human out of the loop:** A review of Bayesian optimization. *Proceedings of the IEEE*, 104, (1), 148-175, doi:10.1109/JPROC.2015.2494218.

# 03 aML

# Best practice examples of aML …

# Recommender Systems



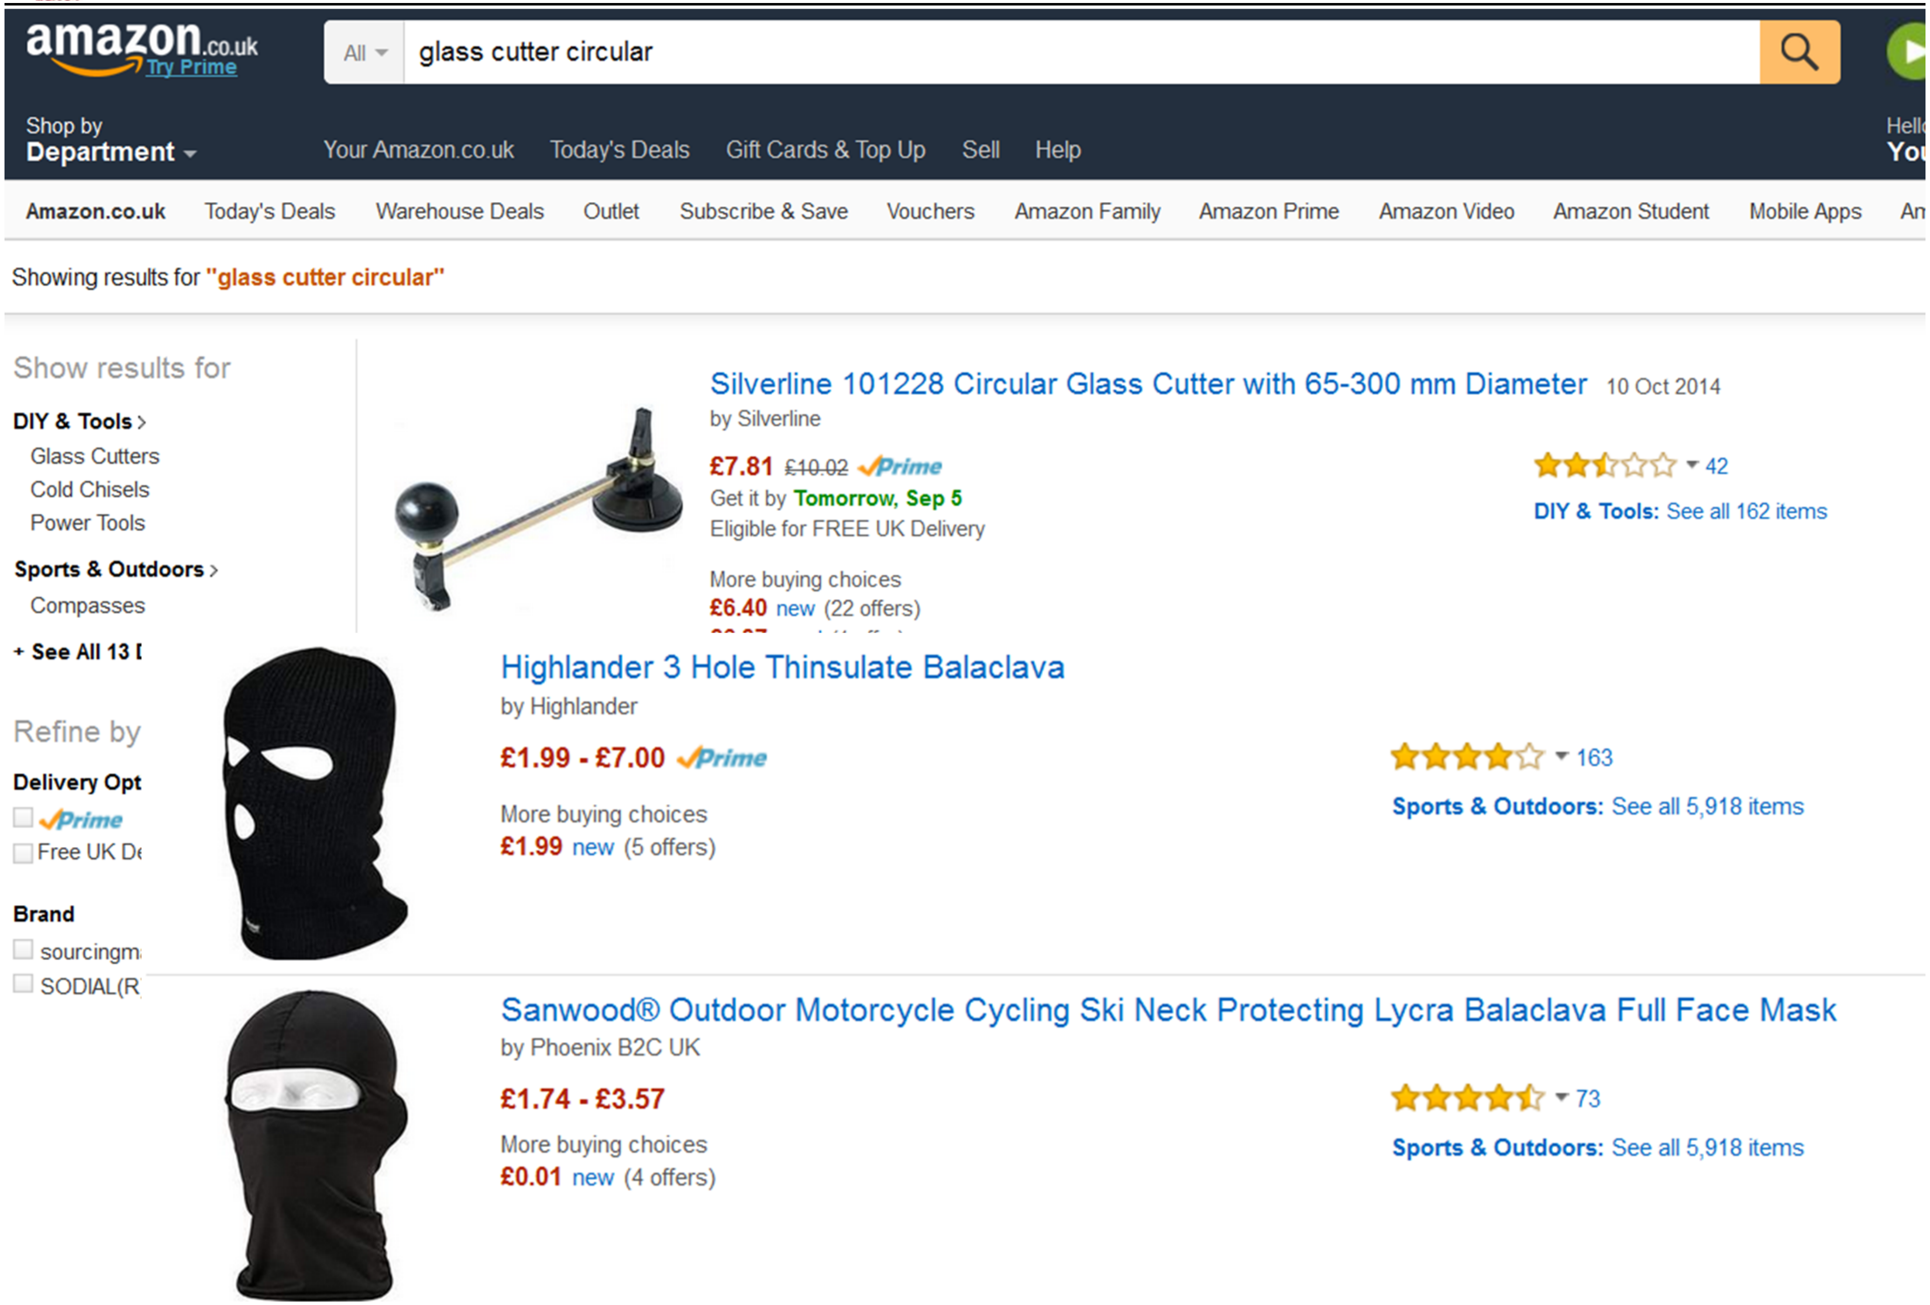

# Fully automatic autonomous vehicles ("Google car")




Guizzo, E. 2011. How google's self-driving car works. IEEE Spectrum Online, 10, 18.

# … an old dream to make it fully automatic

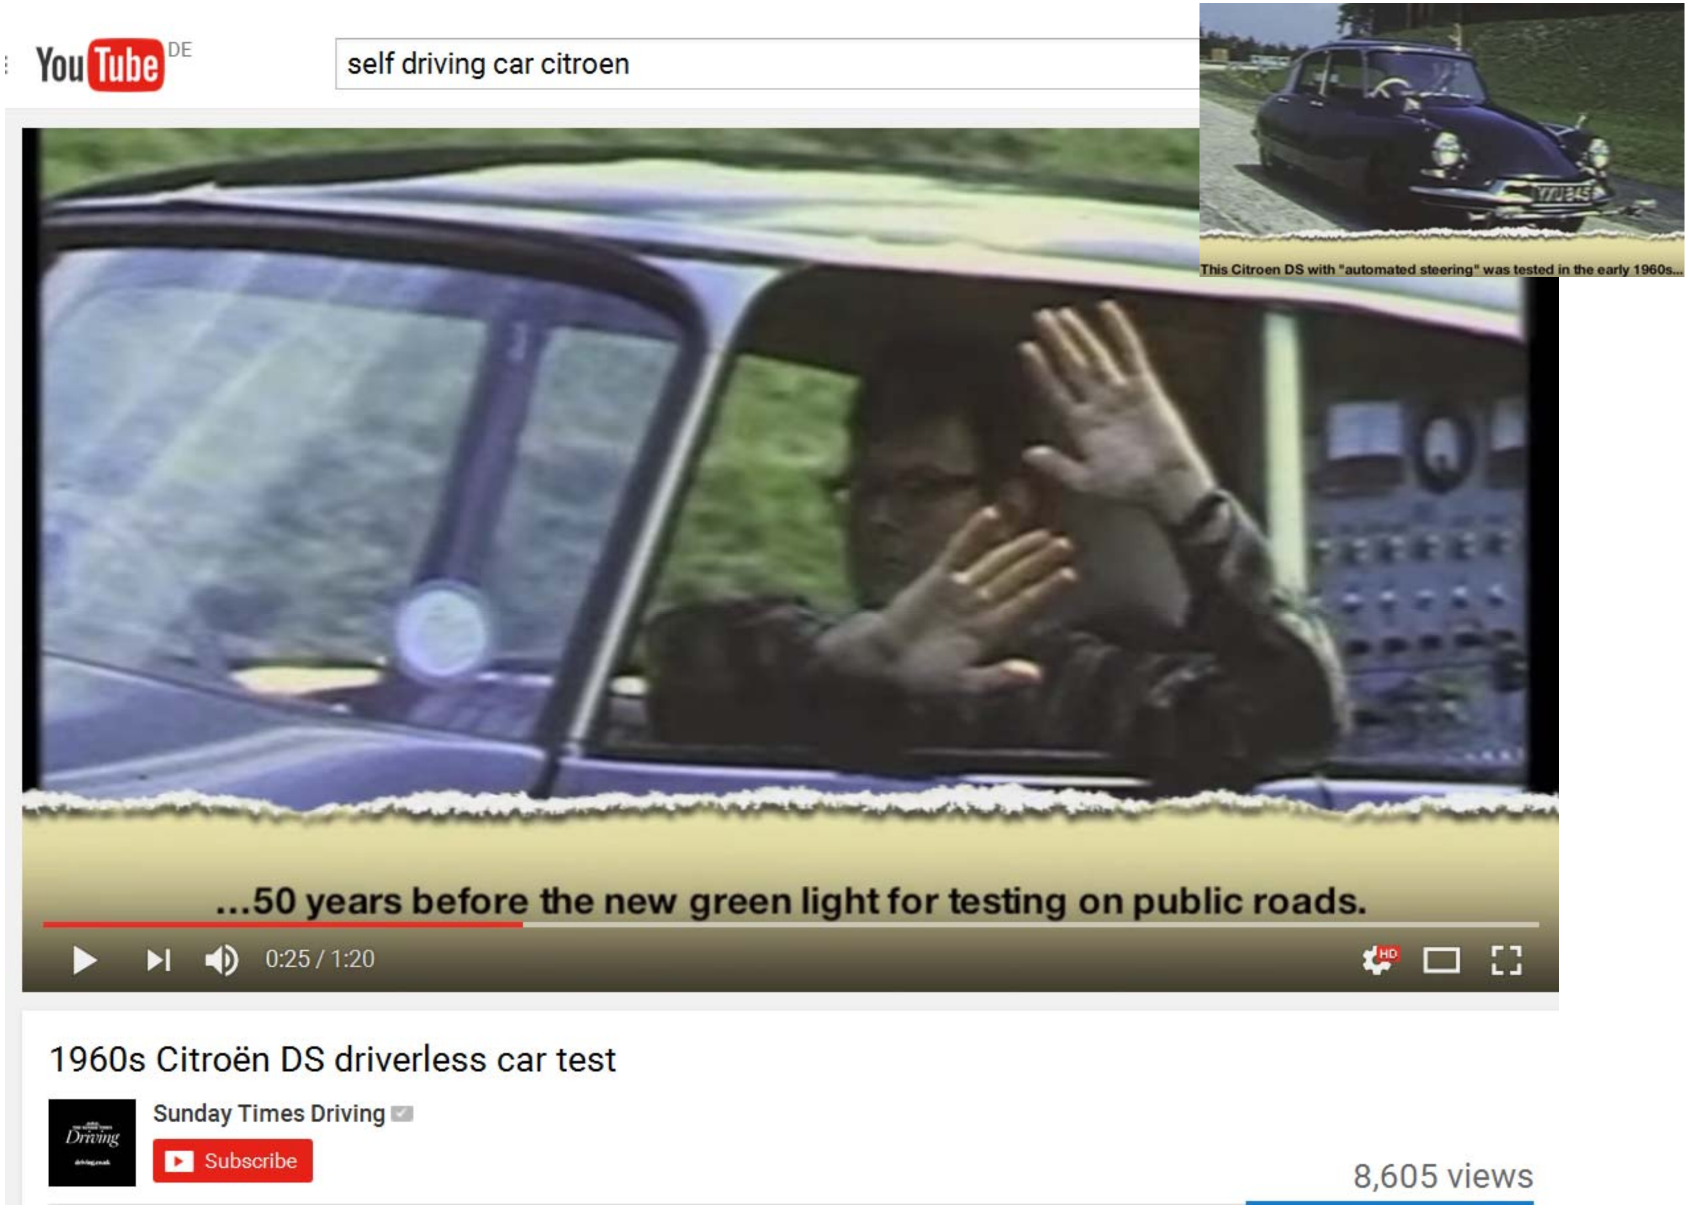

# … and thousands of industrial aML applications …

**Cyber-Physical Systems (CPS):**
***Tight integration of networked computation with physical systems***

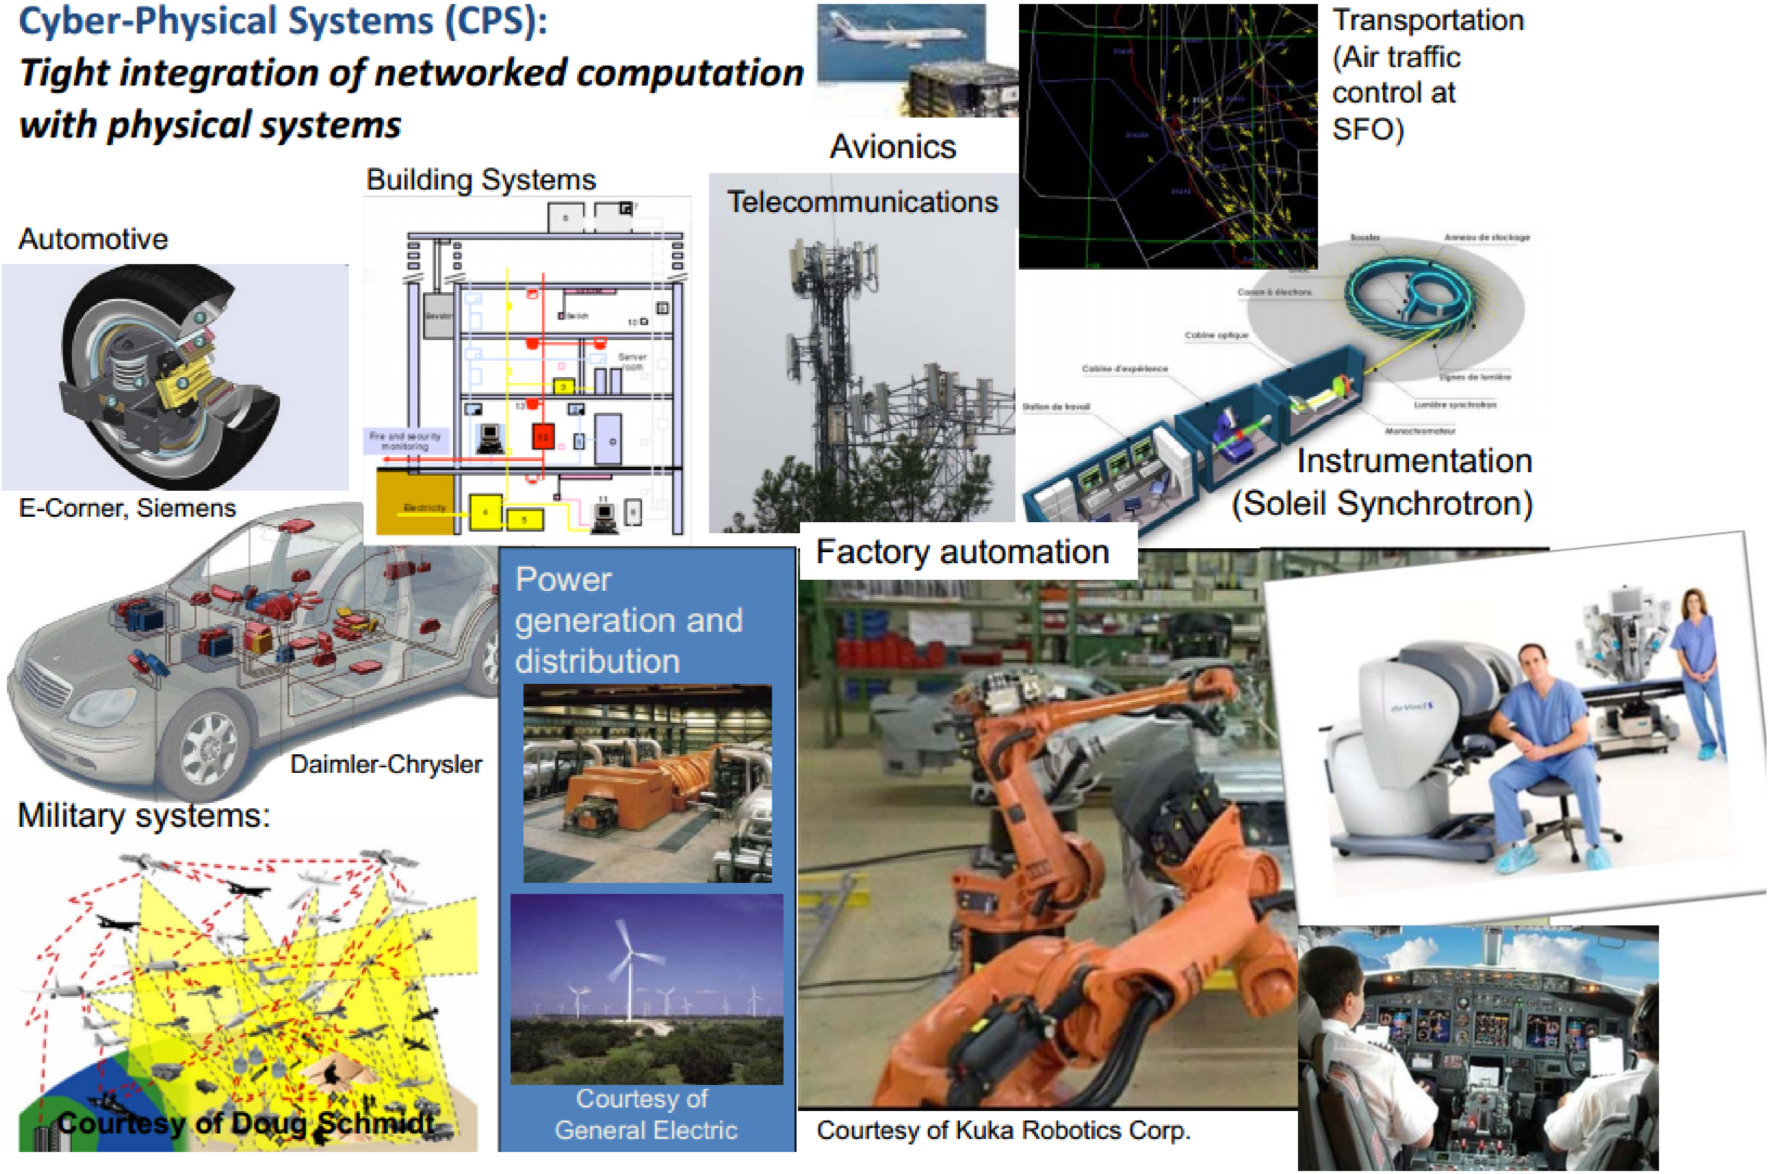


Seshia, S. A., Juniwal, G., Sadigh, D., Donze, A., Li, W., Jensen, J. C., Jin, X., Deshmukh, J., Lee, E. & Sastry, S. 2015. Verification by, for, and of Humans: Formal Methods for Cyber-Physical Systems and Beyond. Illinois ECE Colloquium.

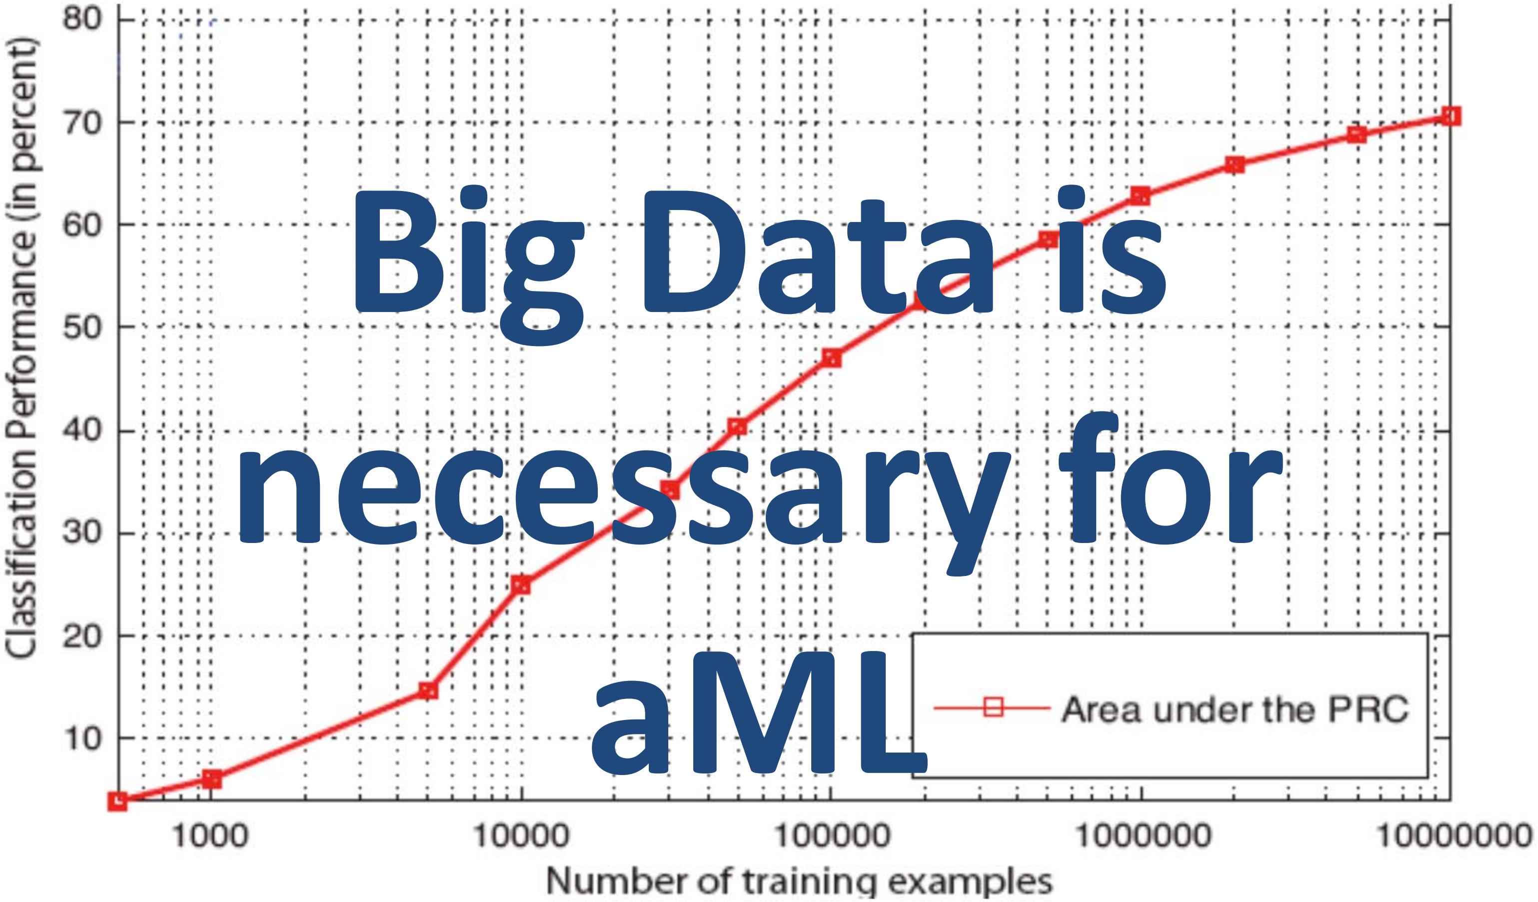


Sonnenburg, S., Rätsch, G., Schäfer, C. & Schölkopf, B. 2006. Large scale multiple kernel learning. Journal of Machine Learning Research, 7, (7), 1531-1565.

# 10 million 200 $x$ 200 px images downloaded from Web

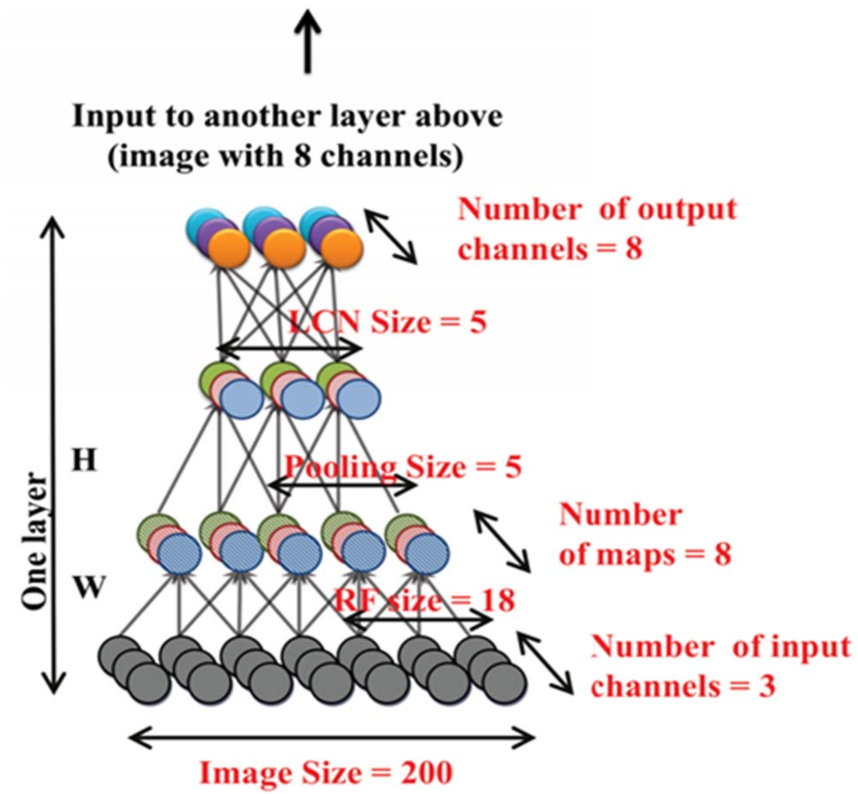


$$x^* = \arg\min_x f(x; W, H), \text{ subject to } ||x||_2 = 1.$$


Le, Q. V., Ranzato, M. A., Monga, R., Devin, M., Chen, K., Corrado, G. S., Dean, J. & Ng, A. Y. 2011. Building high-level features using large scale unsupervised learning. arXiv preprint arXiv:1112.6209.

Le, Q. V. 2013. Building high-level features using large scale unsupervised learning. *IEEE Intl. Conference on Acoustics, Speech and Signal Processing ICASSP.* IEEE. 8595-8598, doi:10.1109/ICASSP.2013.6639343.

# Deep Convolutional Neural Network Pipeline

Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M. & Thrun, S. 2017. Dermatologist-level classification of skin cancer with deep neural networks. Nature, 542, (7639), 115-118, doi:10.1038/nature21056.

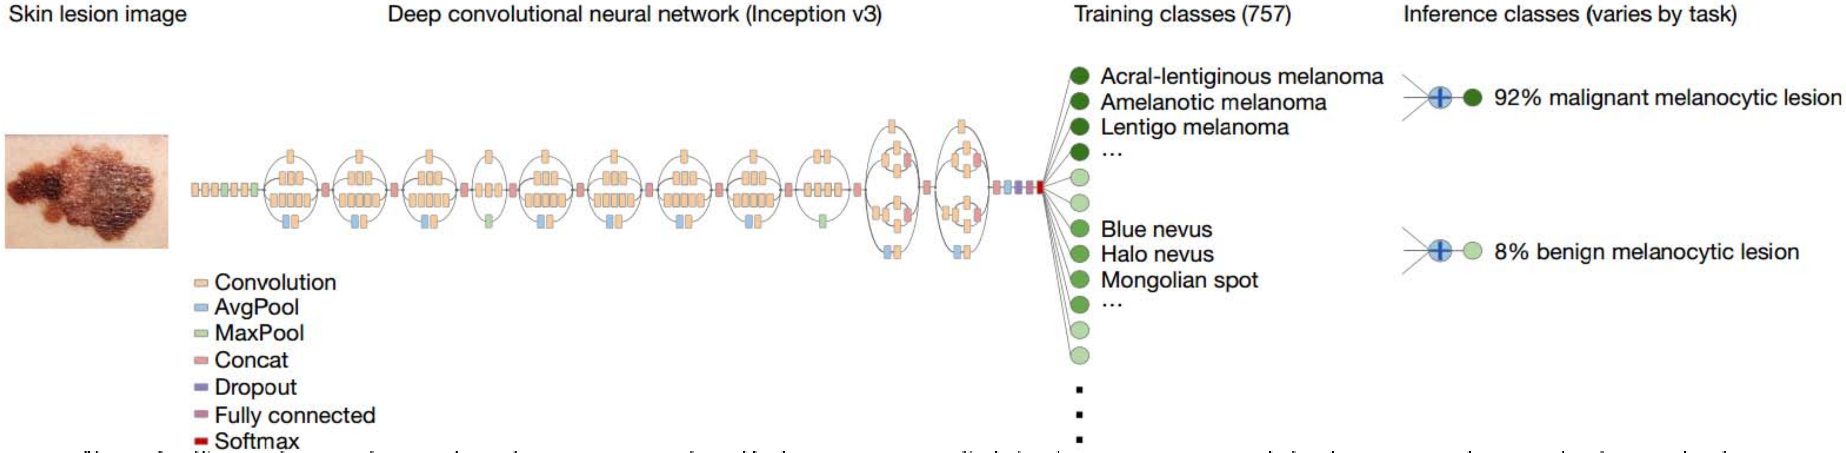


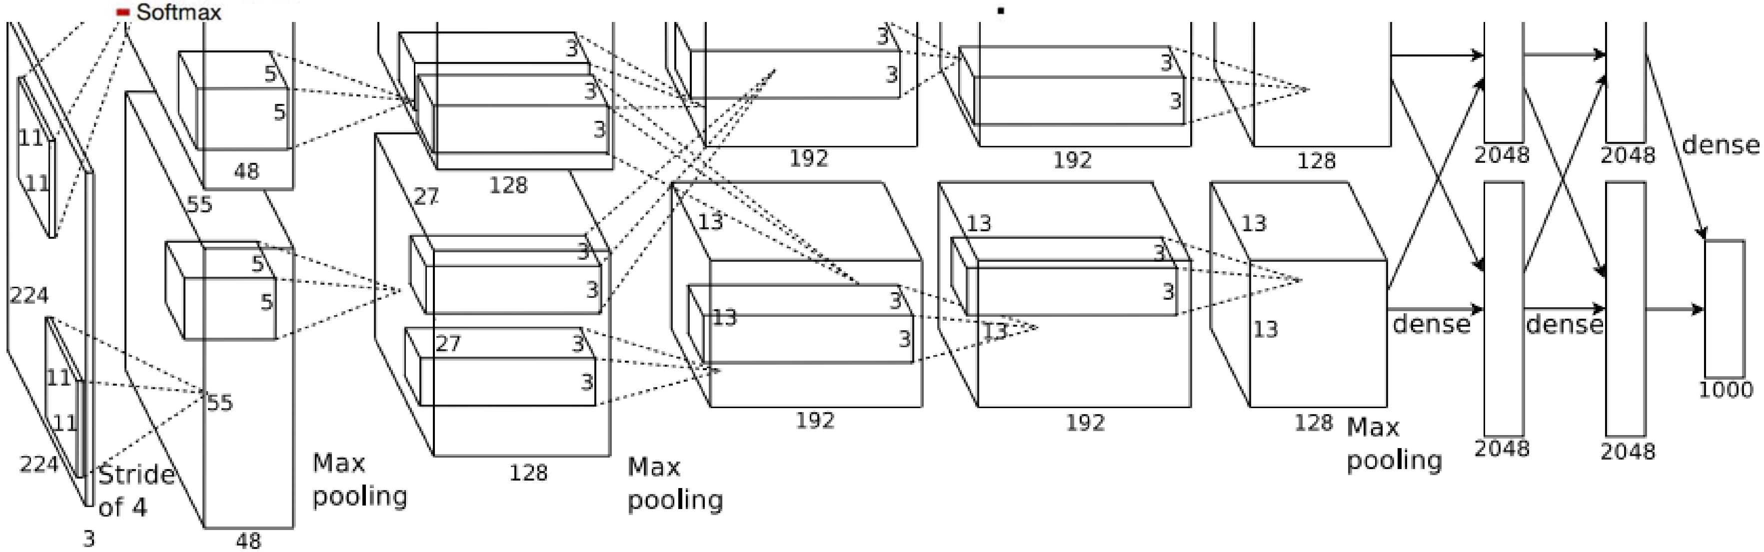


Krizhevsky, A., Sutskever, I. & Hinton, G. E. Imagenet classification with deep convolutional neural networks. In: Pereira, F., Burges, C. J. C., Bottou, L. & Weinberger, K. Q., eds. Advances in neural information processing systems (NIPS 2012), 2012 Lake Tahoe. 1097-1105.

# Limitations of Deep Learning approaches

- Computational resource intensive (supercomps, cloud CPUs, **federated learning,** …)
- Black-Box approaches – lack **transparency,** do not foster trust and acceptance among end-user, legal aspects make “black box” difficult!
- **Non-convex:** difficult to set up, to train, to optimize, needs a lot of expertise, error prone
- Very bad in dealing with **uncertainty**
- Affected by the effect of “catastrophic forgetting”
- **Data intensive, needs often millions of training samples …**

- Sometimes we **do not have “big data”,** where aML-algorithms benefit.
- Sometimes we have
  - **Small amount of data sets**
  - Rare Events – **no training samples**
  - **NP-hard problems, e.g.**
    - Subspace Clustering,
    - k-Anonymization,
    - Protein-Folding, …

# Sometimes we (still) need a human-in-the-loop

# 04 iML

# Definition of iML (Holzinger – 2016)

- iML := algorithms which interact with agents*) and can optimize their learning behaviour through this interaction

***) where the agents can be human**

Holzinger, A. 2016. Interactive Machine Learning for Health Informatics: When do we need the human-in-the-loop? Brain Informatics, 3, (2), 119-131, doi:10.1007/s40708-016-0042-6.

Holzinger, A. 2016. Interactive Machine Learning (iML). Informatik Spektrum, 39, (1), 64-68, doi:10.1007/s00287-015-0941-6.

# Sometimes we need a doctor-in-the-loop

# A group of experts-in-the-loop

# A crowd of people-in-the-loop

# Fits well into the concept of Multi-Agent Systems



# aML: taking the human-out-of-the-loop

A) Unsupervised ML: Algorithm is applied on the raw data and learns fully automatic – Human can check results at the end of the ML-pipeline

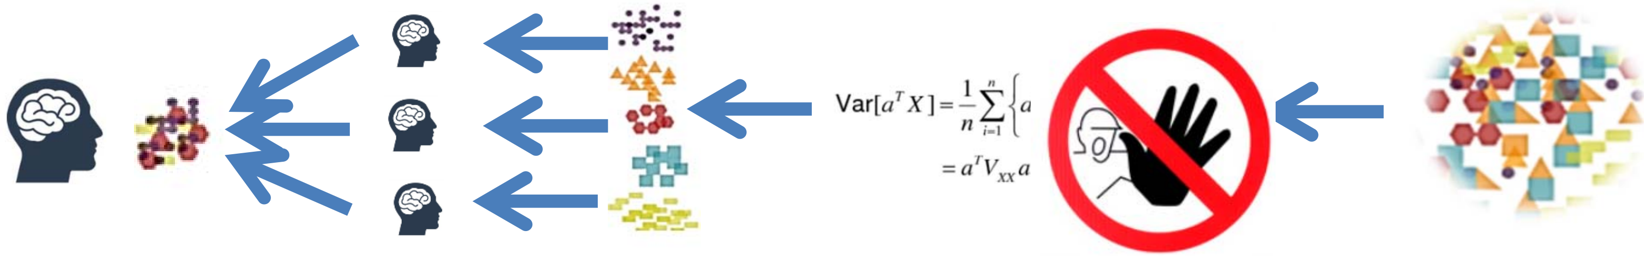

B) Supervised ML: Humans are providing the labels for the training data and/or select features to feed the algorithm to learn – the more samples the better – Human can check results at the end of the ML-pipeline

C) Semi-Supervised Machine Learning: A mixture of A and B – mixing labeled and unlabeled data so that the algorithm can find labels according to a similarity measure to one of the given groups

D) **Interactive Machine Learning:** Human is seen as an agent involved in the actual learning phase, step-by-step influencing measures such as distance, cost functions …

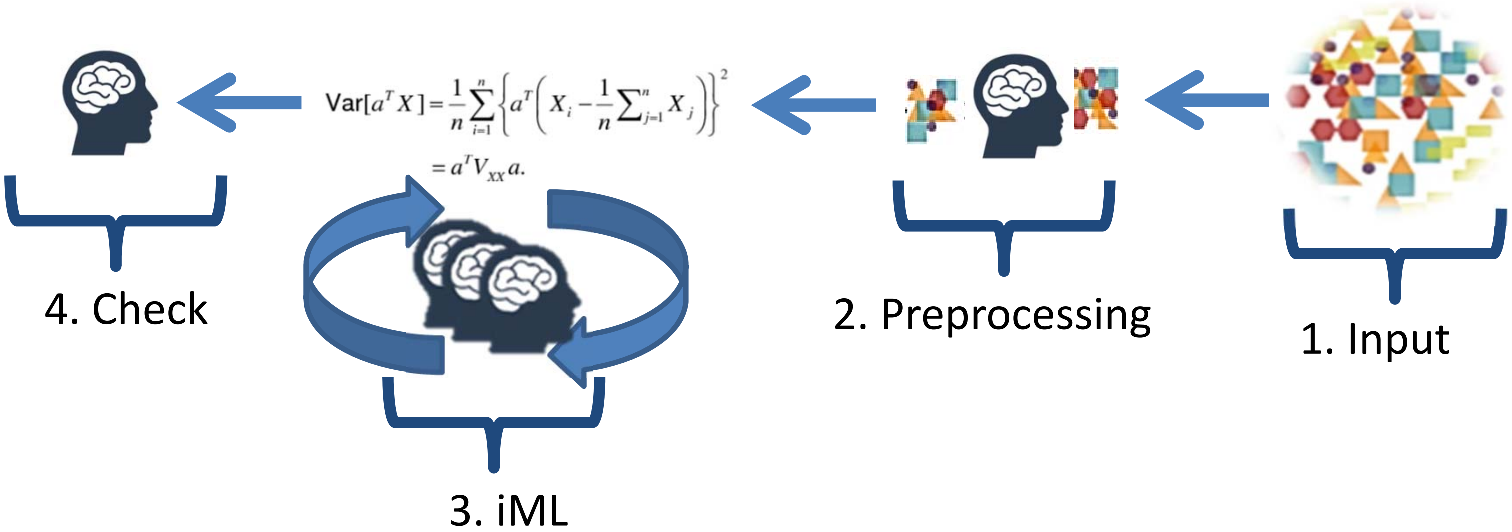


Holzinger, A. 2016. Interactive Machine Learning for Health Informatics: When do we need the human-in-the-loop? Brain Informatics (BRIN), 3, (2), 119-131, doi:10.1007/s40708-016-0042-6.

- **Example 1: Subspace Clustering**
- **Example 2: k-Anonymization**
- **Example 3: Protein Design**

Hund, M., Böhm, D., Sturm, W., Sedlmair, M., Schreck, T., Ullrich, T., Keim, D. A., Majnaric, L. & Holzinger, A. 2016. Visual analytics for concept exploration in subspaces of patient groups: Making sense of complex datasets with the Doctor-in-the-loop. Brain Informatics, 1-15, doi:10.1007/s40708-016-0043-5.

Kieseberg, P., Malle, B., Fruehwirt, P., Weippl, E. & Holzinger, A. 2016. A tamper-proof audit and control system for the doctor in the loop. Brain Informatics, 3, (4), 269–279, doi:10.1007/s40708-016-0046-2.

Lee, S. & Holzinger, A. 2016. Knowledge Discovery from Complex High Dimensional Data. In: Michaelis, S., Piatkowski, N. & Stolpe, M. (eds.) Solving Large Scale Learning Tasks. Challenges and Algorithms, Lecture Notes in Artificial Intelligence LNAI 9580. Springer, pp. 148-167, doi:10.1007/978-3-319-41706-6_7.

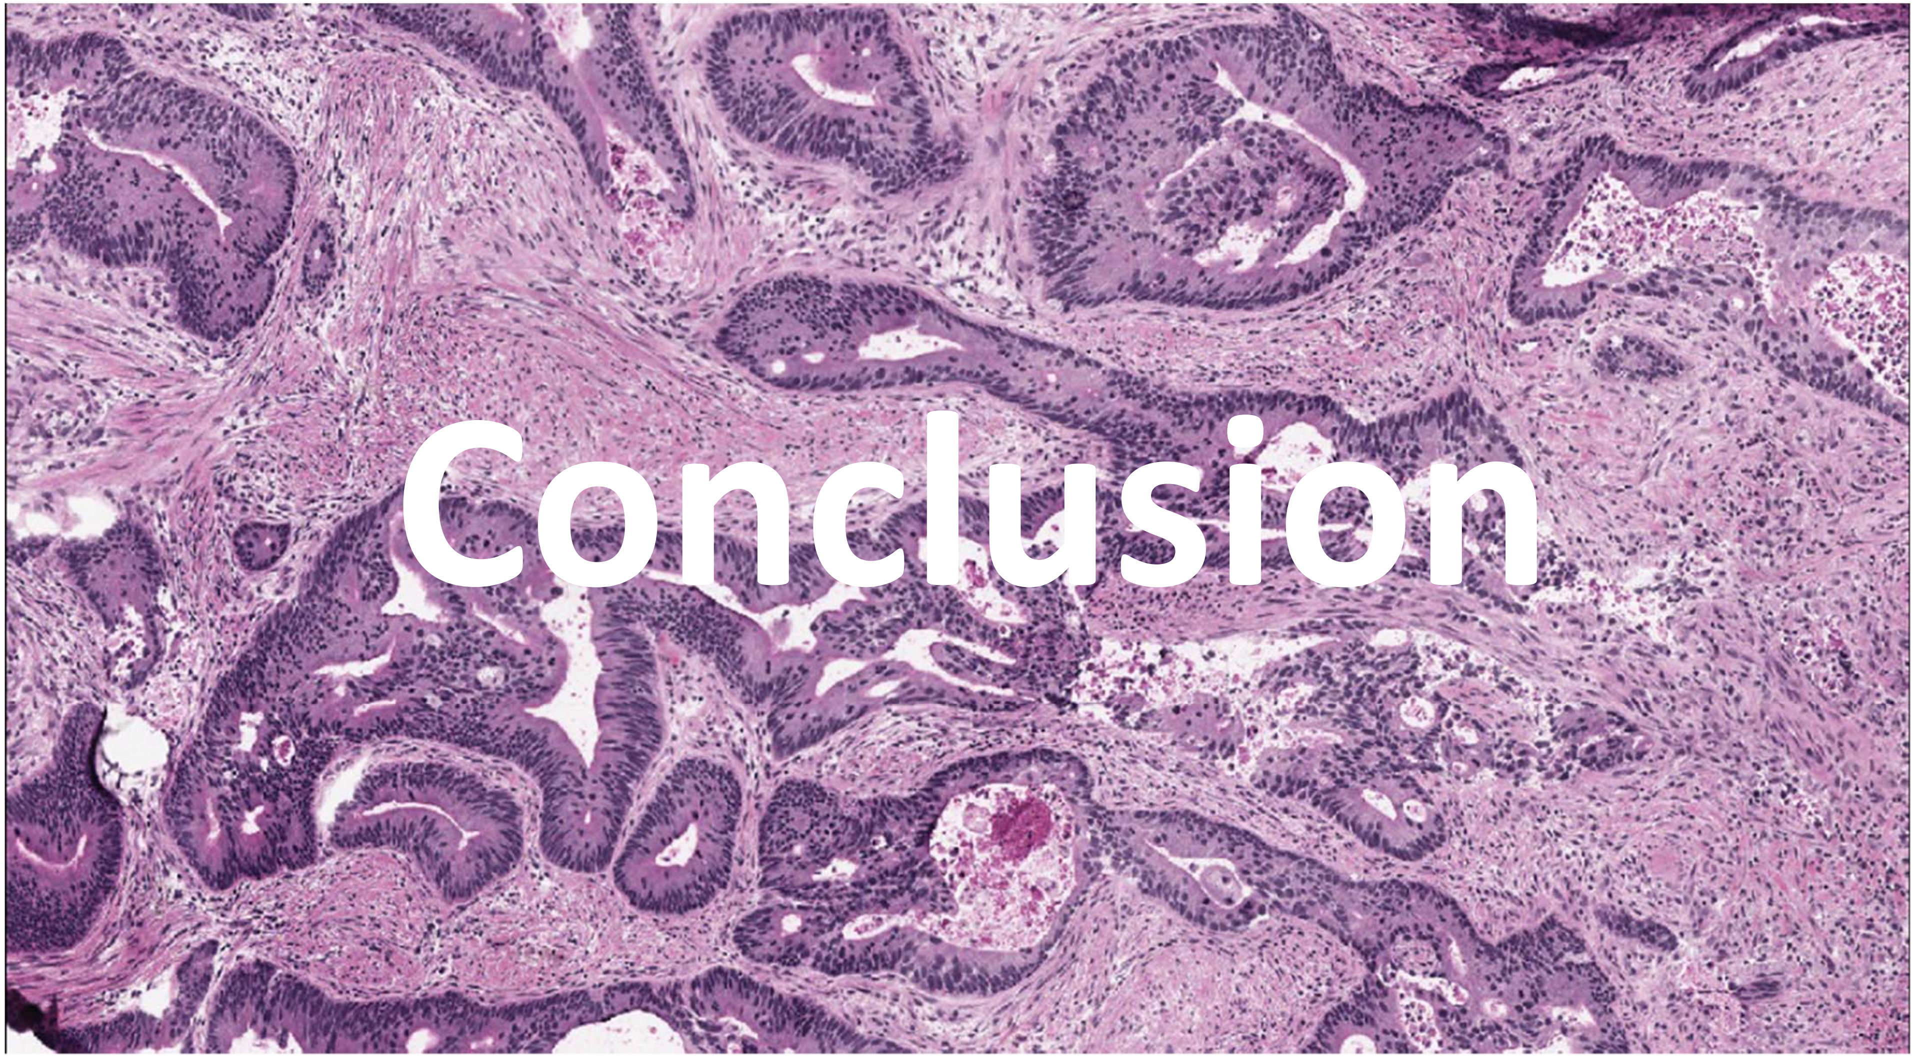
Conclusion

- Computational approaches can find what no human is able to see
- However, still there are many hard problems where a human expert can bring in experience, expertise  knowledge, intuition, …
- **Black box approaches can not explain WHY a decision has been made …**

Holzinger, A., Plass, M., Holzinger, K., Crisan, G.C., Pintea, C.-M. & Palade, V. 2017. A glass-box interactive machine learning approach for solving NP-hard problems with the human-in-the-loop. arXiv:1708.01104.

# Multi-Task Learning …

help to reduce **catastrophic forgetting**

# Transfer learning …

is not easy: learning to perform a task by exploiting knowledge acquired when solving previous tasks:

**A solution to this problem would have major impact to AI research generally and ML specifically!**

# Multi-Agent-Hybrid Systems …

collective intelligence and crowdsourcing
client-side federated machine learning *) – ensures **privacy, data protection, safety & security …**

*) Malle, B., Giuliani, N., Kieseberg, P. & Holzinger, A. 2017. The More the Merrier - Federated Learning from Local Sphere Recommendations. Lecture Notes in Computer Science LNCS 10410 Cham: Springer, pp. 367-374.

# Thank you!