PCC2023

XXI Polish Control Conference

Pursuing Frontiers of Machine Learning: Technology and Society

Witold Pedrycz

Department of Electrical & Computer Engineering
University of Alberta, Edmonton, AB, Canada
and
Systems Research Institute
Polish Academy of Sciences, Warsaw, Poland

wpedrycz@ualberta.ca

Agenda

Machine Learning: main directions, achievements, and challenges

Green Machine Learning

Granular embedding: towards evaluating credibility of models

Transfer learning and federated learning with credibility augmentation

Conclusions and prospects

Machine Learning – bird's eye view

Plethora of learning algorithms processing large amounts of data

Remarkable progress in various areas of applications with Impressive results (natural language processing, computer vision...)

Strategically sound critical areas of applications (autonomous vehicles, healthcare...) with long range impact

Machine Learning – bird's eye view

Plethora of learning algorithms processing large amounts of data

Remarkable progress in various areas of applications with
Impressive results (natural language processing, computer vision...)

Strategically sound critical areas of applications (autonomous vehicles, healthcare...) with long range impact

Enormous computing overhead

Limited interpretability and explainability

Credibility of ML constructs and their solutions

Arising privacy concerns

Brittleness of ML solutions

Society-Oriented Environment of Machine Learning

Creating a holistic view of Machine Learning by understanding society-oriented impact of the discipline and building comprehensive technical solutions

- **Revisiting already existing concepts and methods**
- **Developments of new directions**

Green AI
Explainable AI (XAI)

Society-oriented ML

Granular Computing

Green AI

Enormous computing overhead
Limited interpretability and explainability
Credibility of ML constructs and their solutions
Arising privacy concerns
Brittleness of ML solutions

XAI

Transfer Learning

Federated Learning

Credibility assessment

From ML to society oriented ML

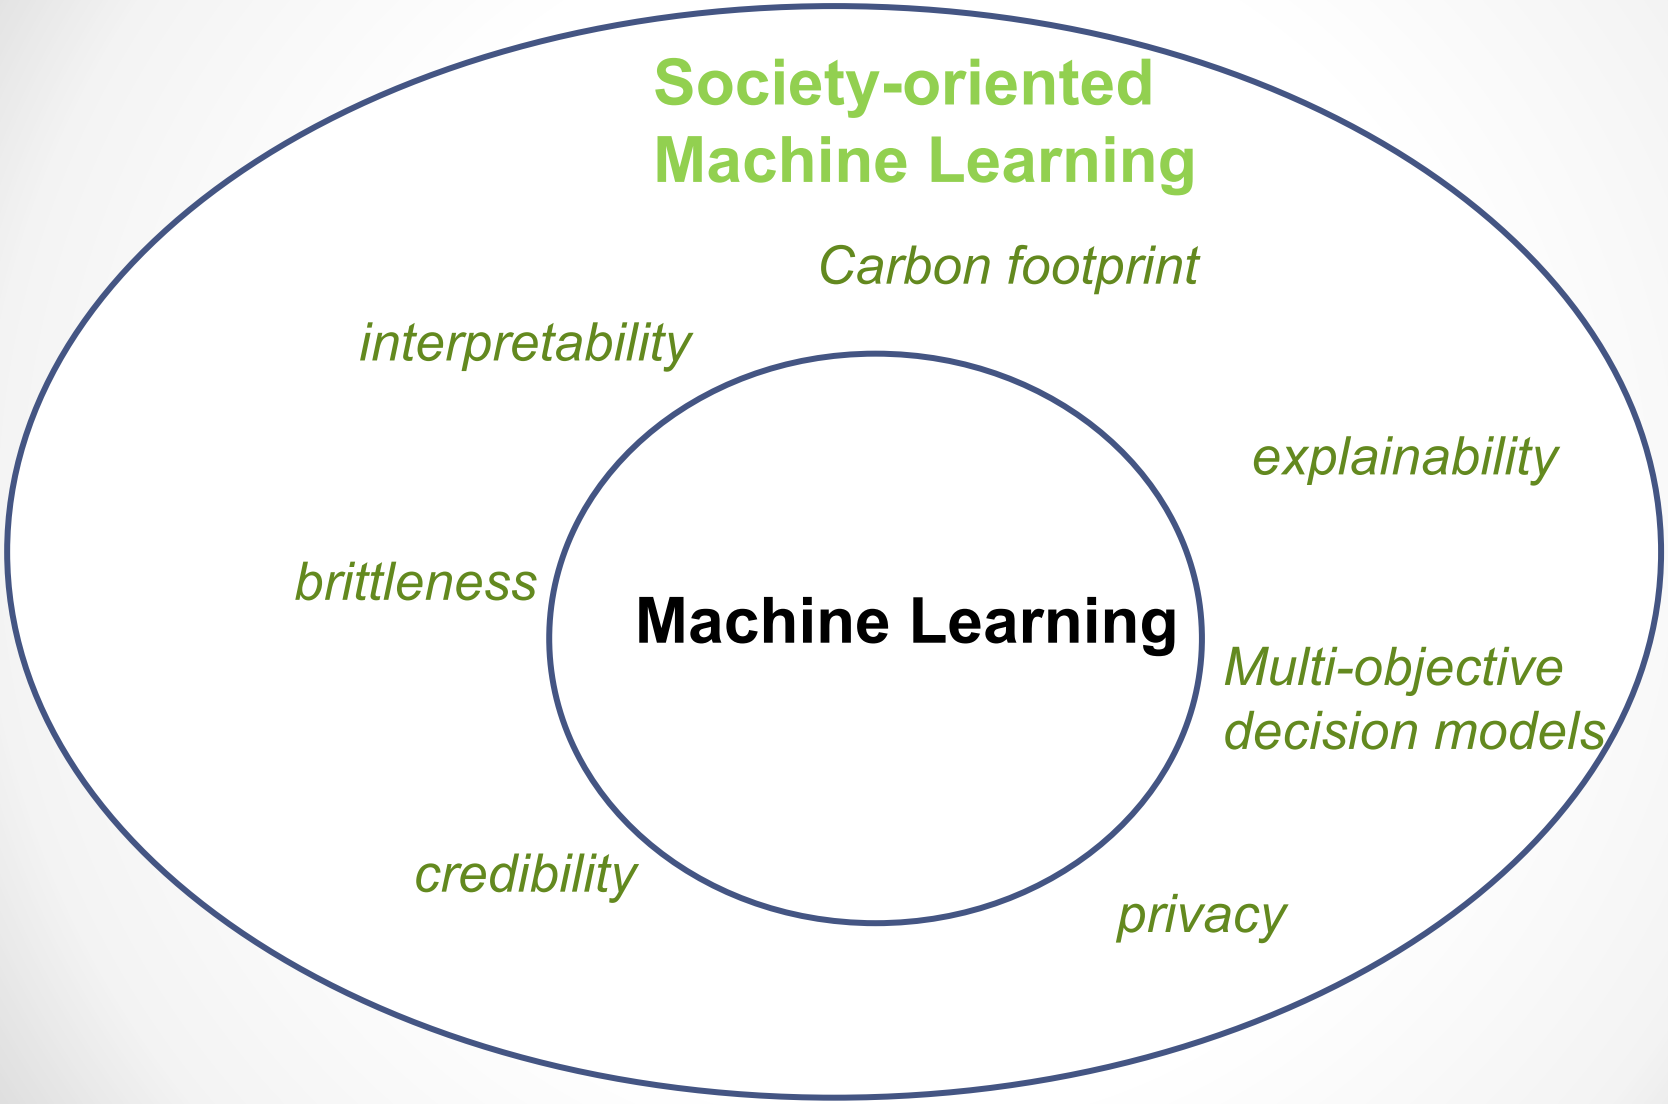

Green Machine Learning

Green AI (ML) and
Green Machine Learning

Computing to realize deep learning doubles every few months

from 2012 to 2018- 300,000 increase of required computing

Huge number of parameters (connections) to learn

R. Schwartz, J. Dodge, et al., Green AI, 2019

Green AI (ML) and Green Machine Learning

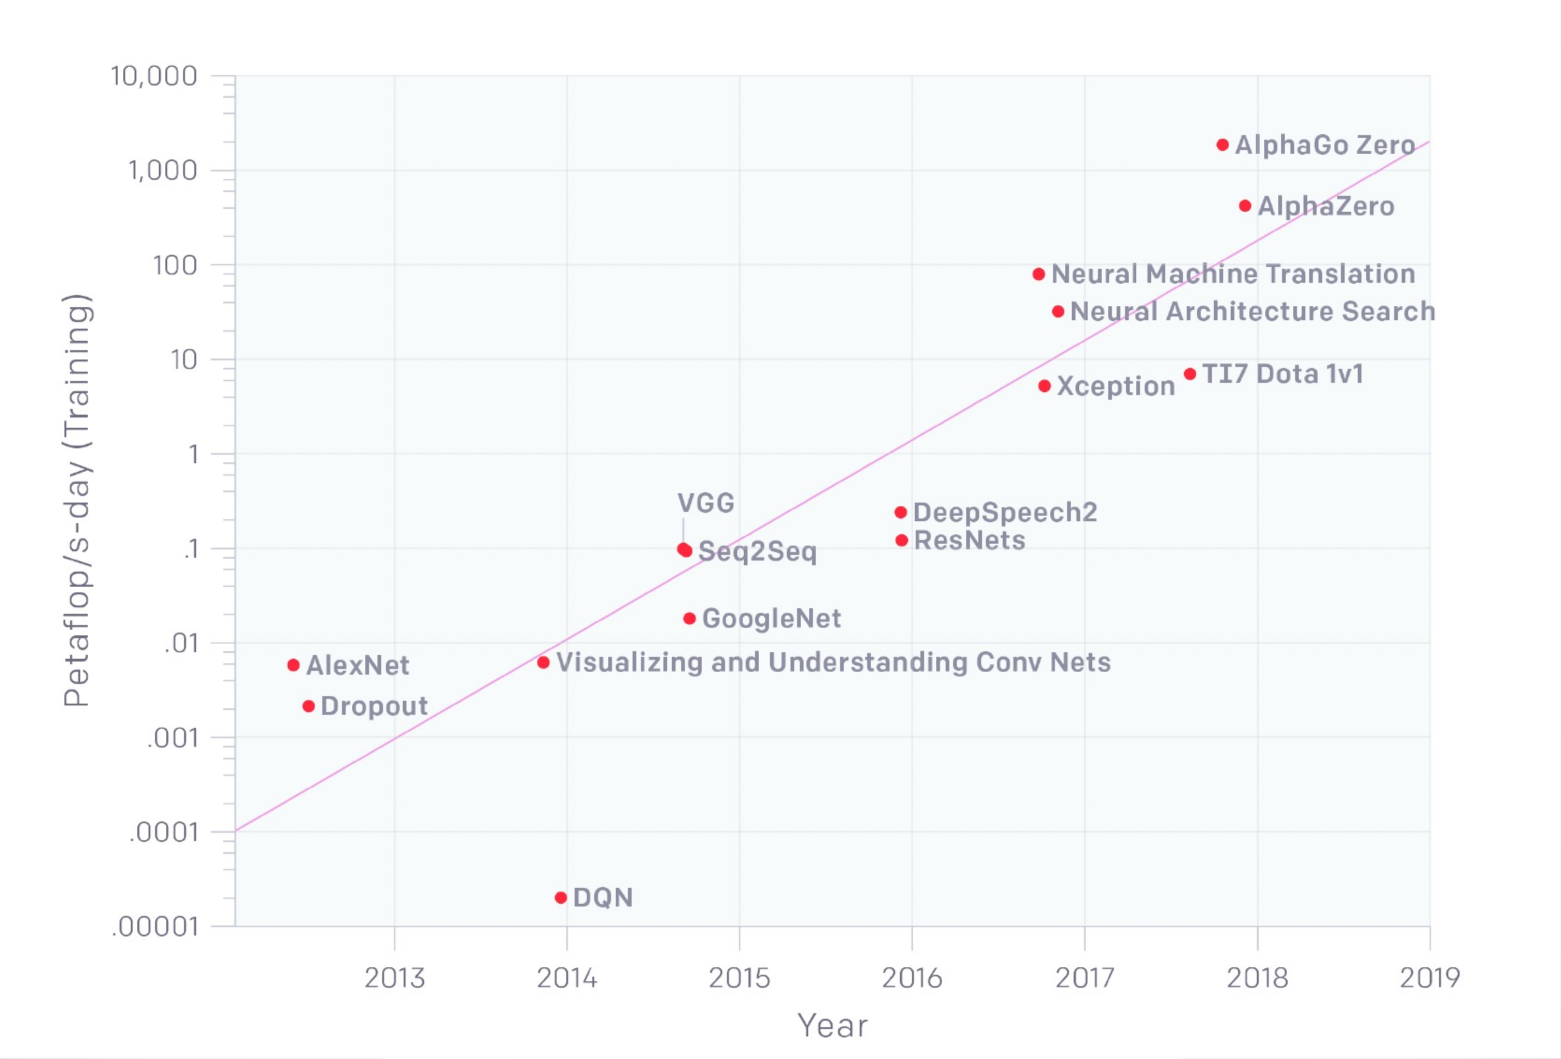

ARTIFICIAL INTELLIGENCE

An AI that writes convincing prose risks mass-producing fake news

Fed with billions of words, this algorithm creates convincing articles and shows how AI could be used to fool people on a mass scale.

It takes a lot of energy for machines to learn – here's why AI is so power-hungry

MIT Technology Review

Sign in

Subscribe

ARTIFICIAL INTELLIGENCE

Training a single AI model can emit as much carbon as five cars in their lifetimes

Deep learning has a terrible carbon footprint.

Selected numbers

Natural Language Processing

GPT-2, has 1.5 billion weights in its network.
GPT-3, has 175 billion weights.

Carbon footprint

training an AI model generates as much carbon emissions as it takes to build and drive five cars over their lifetimes training.

training BERT once has the carbon footprint of a passenger flying a round trip between New York and San Francisco. However, by searching using different structures – that is, by training the algorithm multiple times on the data with slightly different numbers of neurons, connections and other parameters – the cost became the equivalent of 315 passengers, or an entire 747 jet.

GPT 4

Generative Pretrained Transformer

Chat GPT 2: 1.5 billion

Chat GPT 3: 175 billion parameters, **936 MWh**
Household per year: 10,632 kWh
953.7 lbs CO_2 per 1 MWh for delivered electricity

Chat GPT 4: 175 billion parameters, 1 trillion?
Language model: Text generation,
language translation,
language generation,
Automated content generation...

ML constructs:
Energy consumption and carbon footprint

	Date of original paper	Energy consumption (kWh)	Carbon footprint (lbs of CO2e)
Transformer (65M parameters)	Jun, 2017	27	26
Transformer (213M parameters)	Jun, 2017	201	192
ELMo	Feb, 2018	275	262
BERT (110M parameters)	Oct, 2018	1,507	1,438
Transformer (213M parameters) w/ neural architecture search	Jan, 2019	656,347	626,155

MIT Technology Review

MIT Technology Review

Common carbon footprint benchmarks

in lbs of CO2 equivalent

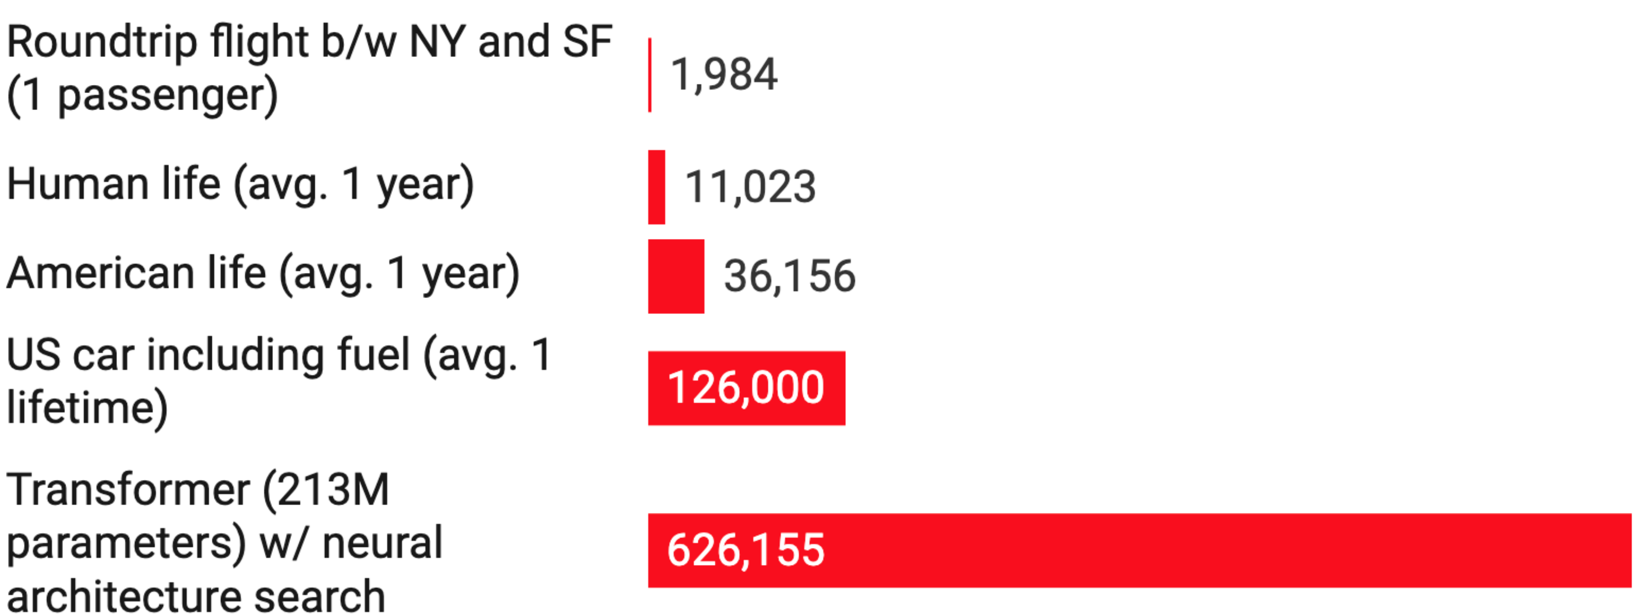

Green AI (ML) and Green Machine Learning

Dominant direction:

Buying “stronger” results (**accuracy**) by engaging massive computing power; limited return on investment?

Green AI (ML) and Green Machine Learning

Striving for efficiency of ML constructs:

computing overhead versus improved accuracy

balance of **efficiency**

Carbon emission

Electricity usage

Elapsed real time

Floating point operations

Towards Green ML

ML Model

Transfer Learning

Knowledge Distillation

Federated Learning

Granular ML Model

Design with
granular regularization

Performance evaluation
with granular models

Information granules and Granular Computing

Information granules: pieces of knowledge resulting as an abstraction of data, exhibiting well-defined semantics and forming functional modules in further interpretable system modeling (granular models).

Formal frameworks of sets (intervals), fuzzy sets, rough sets,...
z-numbers (Zadeh, 2011)

Granular Computing: knowledge-based environment supporting the design and processing of information granules

Granular Computing for Machine Learning models

Designing of ML models at a suitable level of abstraction

Coping with uncertain (granular) experimental data

Delivering Interpretability and explainability mechanisms (e.g, rules)

Quantifying credibility of the model and its results

Explainable ML

Interpretability (1)

Interpretability: a notion

Results that are easily comprehended by the user producing semantically sound and actionable findings.

Interpretability (2)

Numbers versus information granules

temperature is 25C

no context – space not specified

temperature is *high*

Explainability

Explainability

modeling faculties to:

produce knowledge about relationship existing in data/models and help explain and audit prediction/classification results in response to issues of regulatory or fairness nature

support “what-if” analysis.

support traceability of the reasoning (inference) process.

why did the model produce a particular prediction?

why weren’t other decisions made?

Interpretability and explainability

Required levels of abstraction (details)
pivotal role of information granularity

Flexibility

Actionability

Orientation on user/recipient of ML models

Explainable AI (XAI)

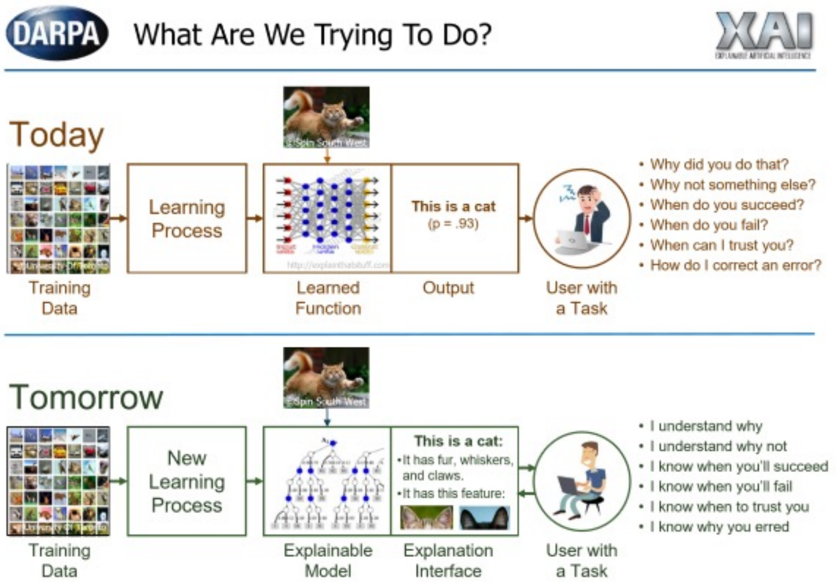

Explainable models: understand, trust, manage produced results

From: D. Gunning, DARPA, 2017

Learning, accuracy, and interpretability capabilities

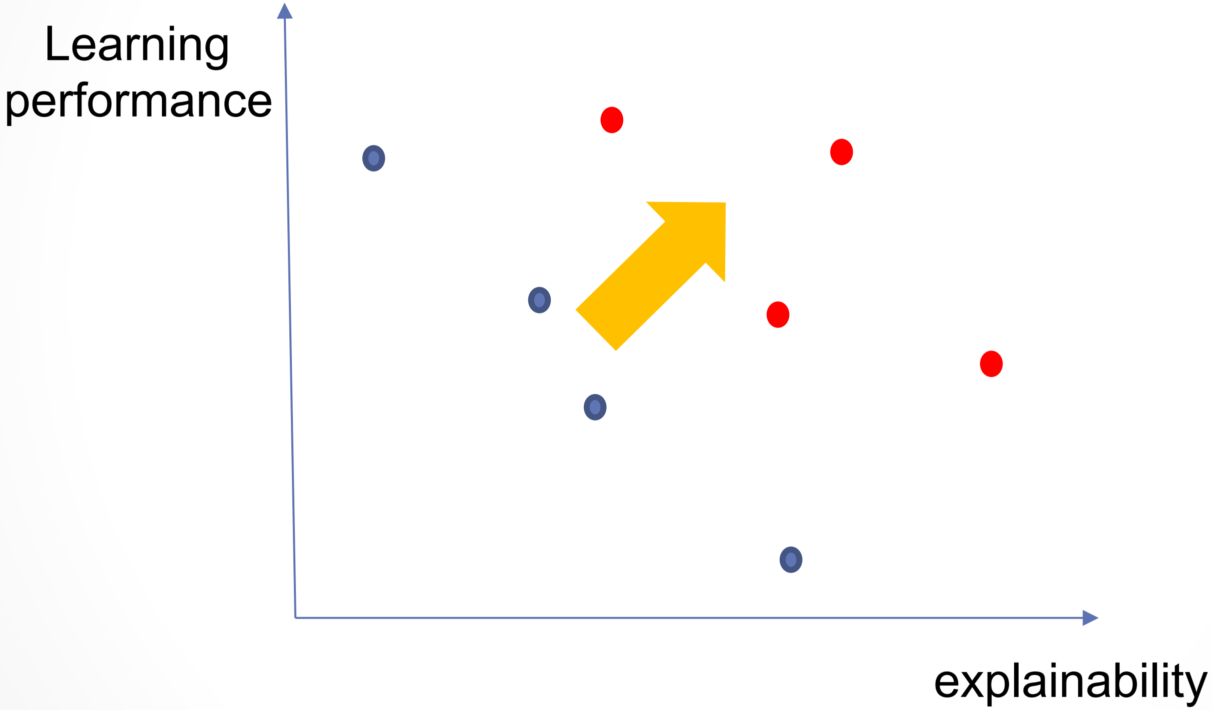

Credibility of ML models

Credibility of Machine Learning Models

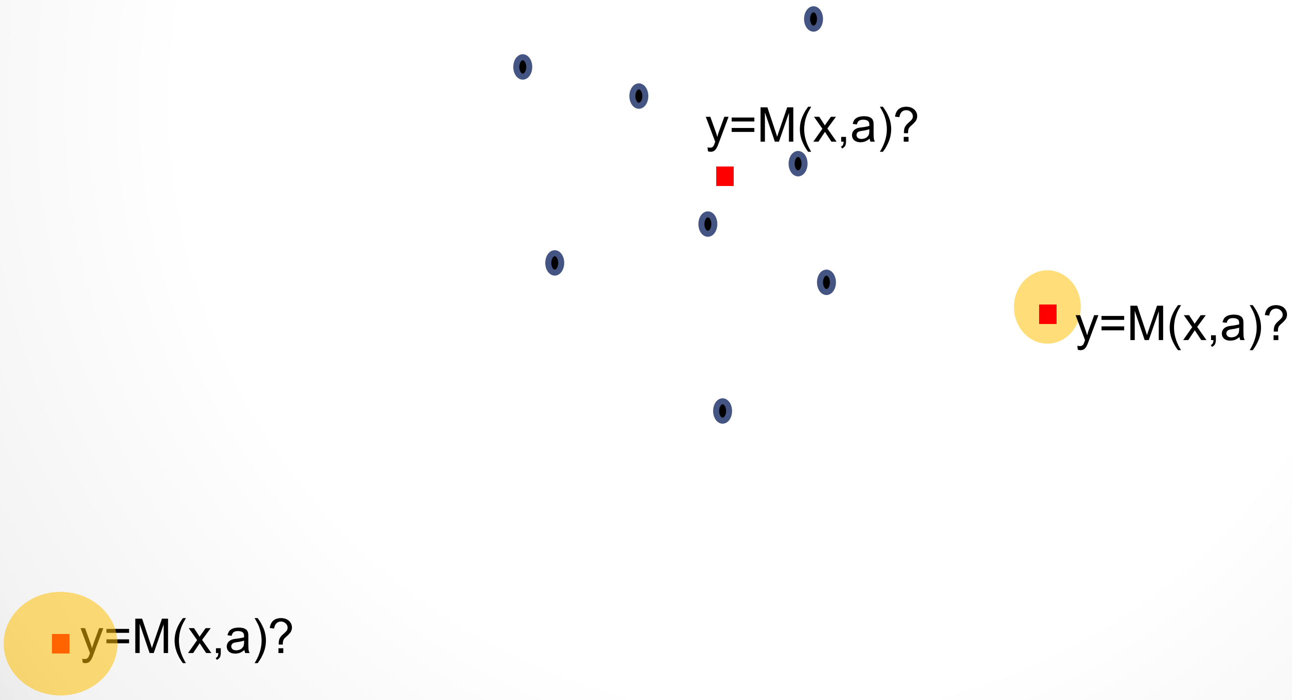

Credibility of Machine Learning Models

New $\boldsymbol{x}$, result $M(\boldsymbol{x}; \boldsymbol{a}_{opt})$

How credible is the result ?

How much confidence could be associated with the result?

Could any action /decision be taken on a basis of obtained result; self-awareness mechanism

Credibility of the model: Granular augmentation of results

Raising and quantifying **awareness** about quality of results

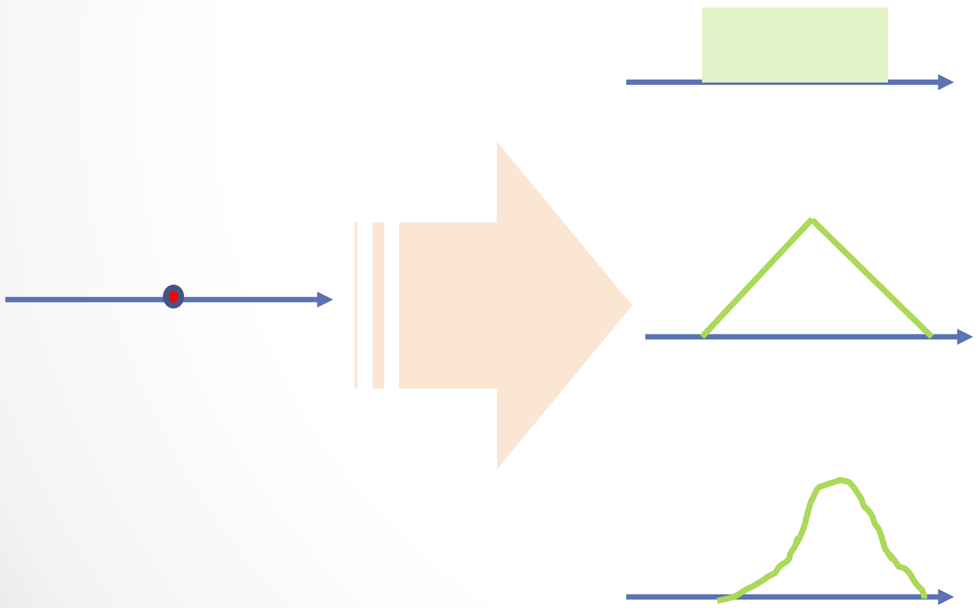

Interval information granule
(prediction interval)

Fuzzy set
information granules

Probability information granule
(probability function)

From numeric results to information granules

Confidence interval (probabilistic information granule)

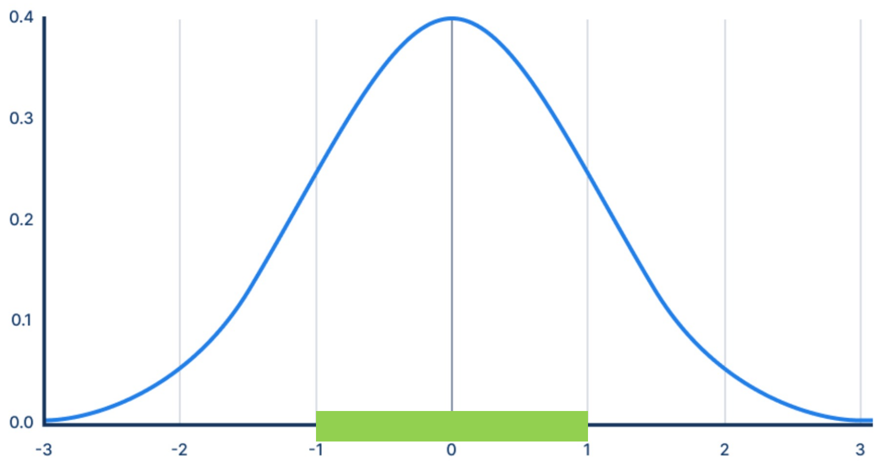

Probability of coverage α=0.05,. 0.01

$$P(x \in A) = 1 - \alpha$$

From numeric to granular models

Linear regression

confidence and prediction intervals

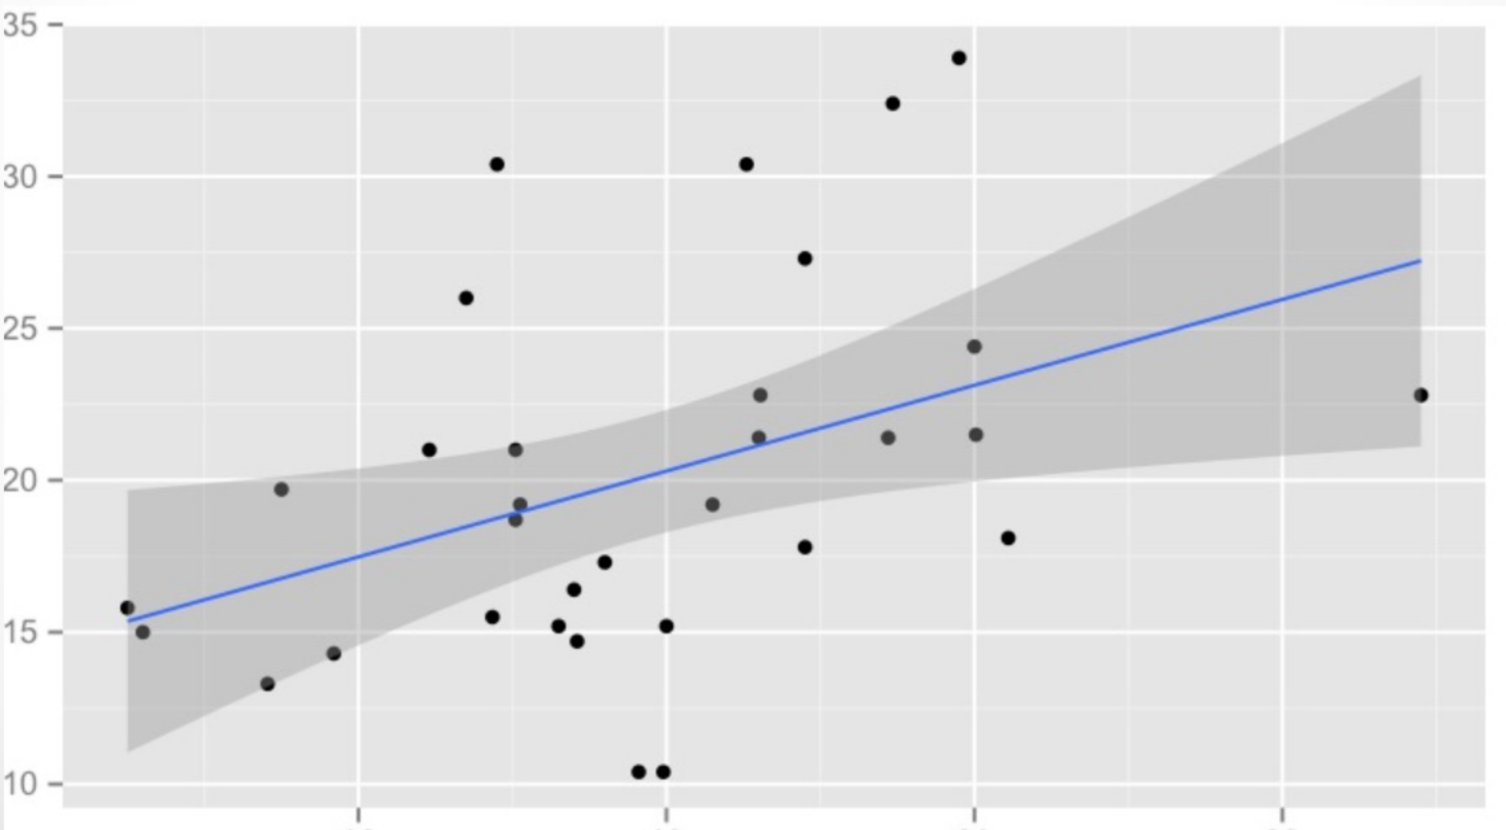

From models to granular models: design asset of information granularity (ε)

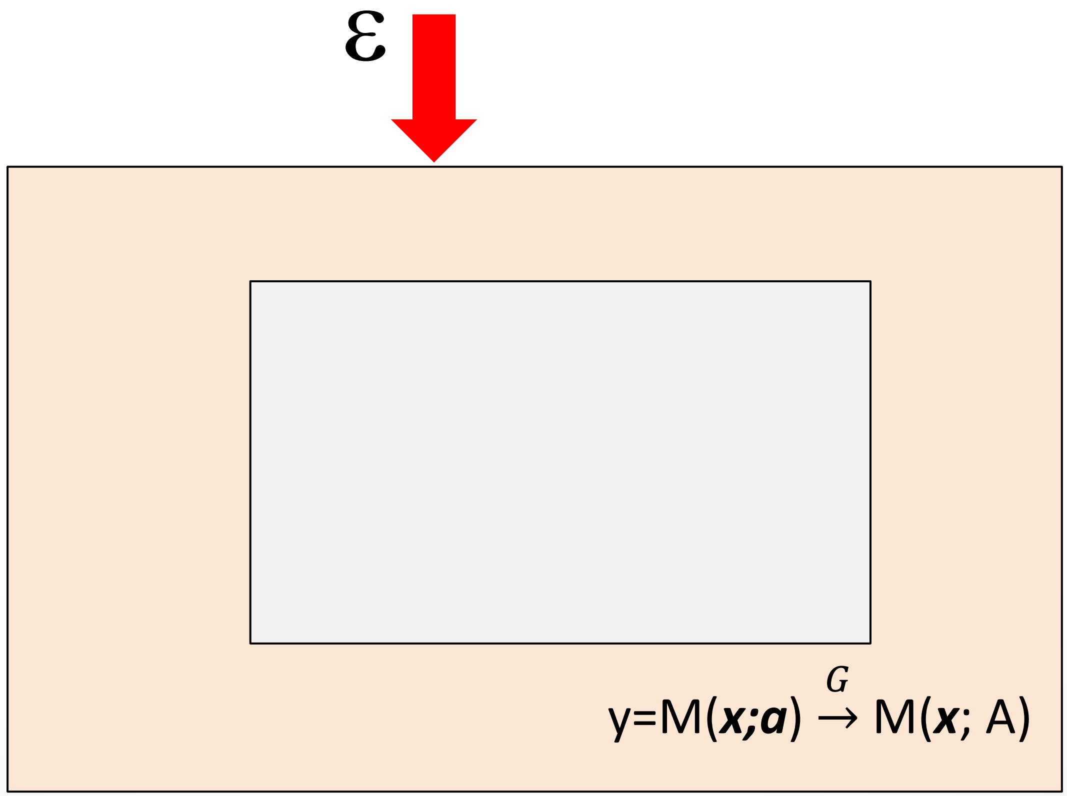

Coverage and specificity

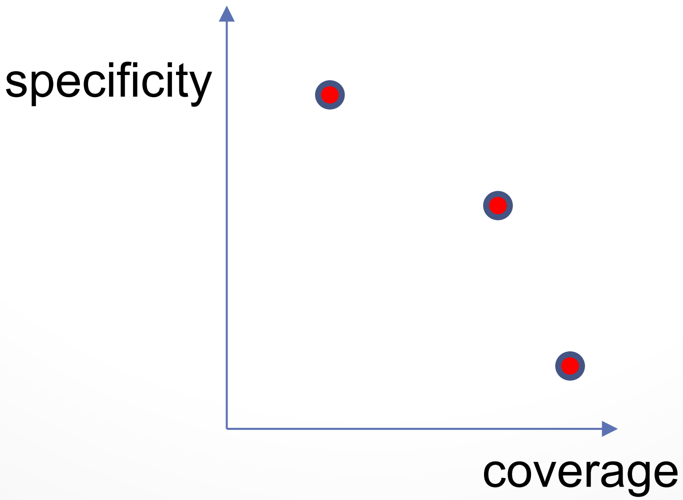

Granular Embedding

$$M \xrightarrow{G} G(M)$$

Granular elevation of parameters

$$y = M(\boldsymbol{x};\boldsymbol{a}) \xrightarrow{G} Y = M(\boldsymbol{x};G(\boldsymbol{a})) = M(\boldsymbol{x};A)$$

Granular elevation of parameters- level of information granularity (ε)

$$y=M(\boldsymbol{x};\boldsymbol{a}) \xrightarrow{G} Y=M(\boldsymbol{x};G(\boldsymbol{a}))= M(\boldsymbol{x}; A)$$

Transformation #1:

$$a \xrightarrow{\epsilon} [\min(a_i(1+\varepsilon),a_i(1-\varepsilon)), \max(a_i(1+\varepsilon),a_i(1-\varepsilon))], \ \varepsilon \in [0,1]$$

Transformation #2:

$$a \xrightarrow{\epsilon} [\min(a_i(1+\varepsilon),a_i/(1+\varepsilon)), \max(a_i(1+\varepsilon),a_i/(1-\varepsilon))], \ \varepsilon \geq 0$$

Performance of granular model

$$cov = \frac{1}{N}\sum_{k=1}^{N} incl(target_k, Y_k)$$

incl(b, B)= $\begin{cases} 1 \; if \; b \in B \\ 0, otherwise \end{cases}$

$$sp = \frac{1}{N}\sum_{k=1}^{N} g(length(Y_k))$$

g-decreasing function of length of Y_k

ε= arg max$_\varepsilon$(cov*sp)

Optimization protocol: level of information granularity

The same level of information granularity ε across all parameters

$$\varepsilon = \arg\max_{\varepsilon}(cov*sp)$$

Individual levels of information granularity associated with parameters $\varepsilon_1, \varepsilon_2, \ldots, \varepsilon_p$, p-number of parameters

$$(\varepsilon_1, \varepsilon_2, \ldots \varepsilon_p) = \arg\max_{\varepsilon}(cov*sp)$$

Data privacy

Federated Learning

Credibility of ML models and results

Granular Computing

Federated Learning

Building a holistic model in the presence of distributed and non-shared data (data islands):

*requirements of privacy and security

*unreliable and limited communication links

*legal requirements (General Protection Regulations; China Security Law of PRC, etc.)

Federated Learning: Paradigm shift

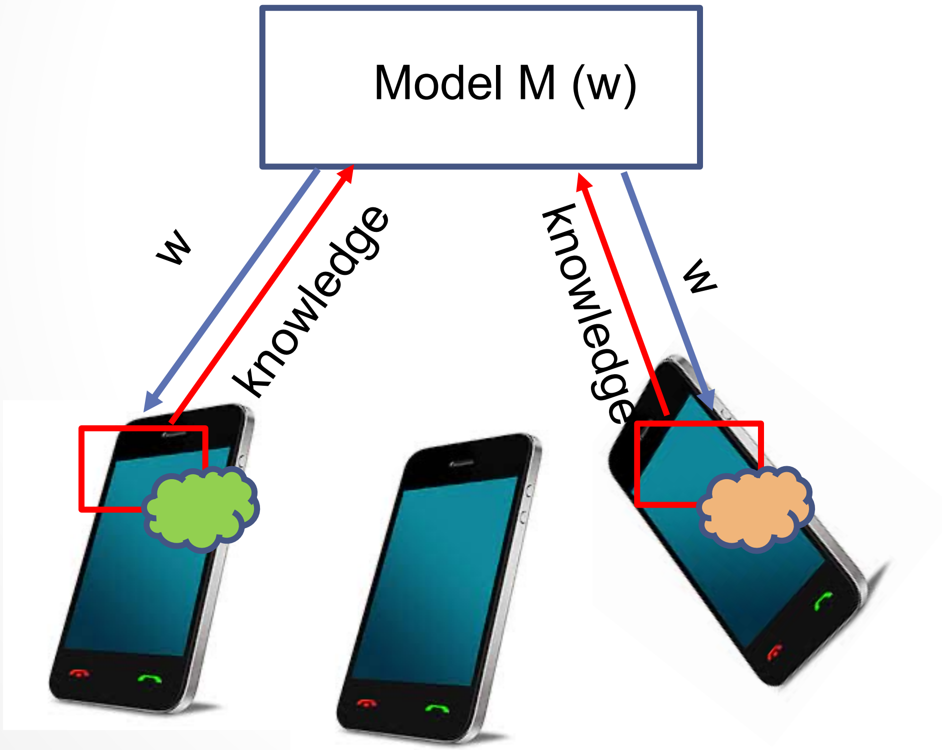

server

Clients (1, 2,...,p)

Federated learning: applications

Education
Deep knowledge tracking system

Healthcare
Privacy-preserving platform
Decentralized optimization framework
Prediction mortality, delivery prediction

Internet of Things (IoT)
Data sharing architecture
intelligent resource management

Smart Transportation
Protecting privacy in traffic flow prediction
Traffic collision avoidance
Optimization of vehicular communications

Federated Learning: Paradigm shift

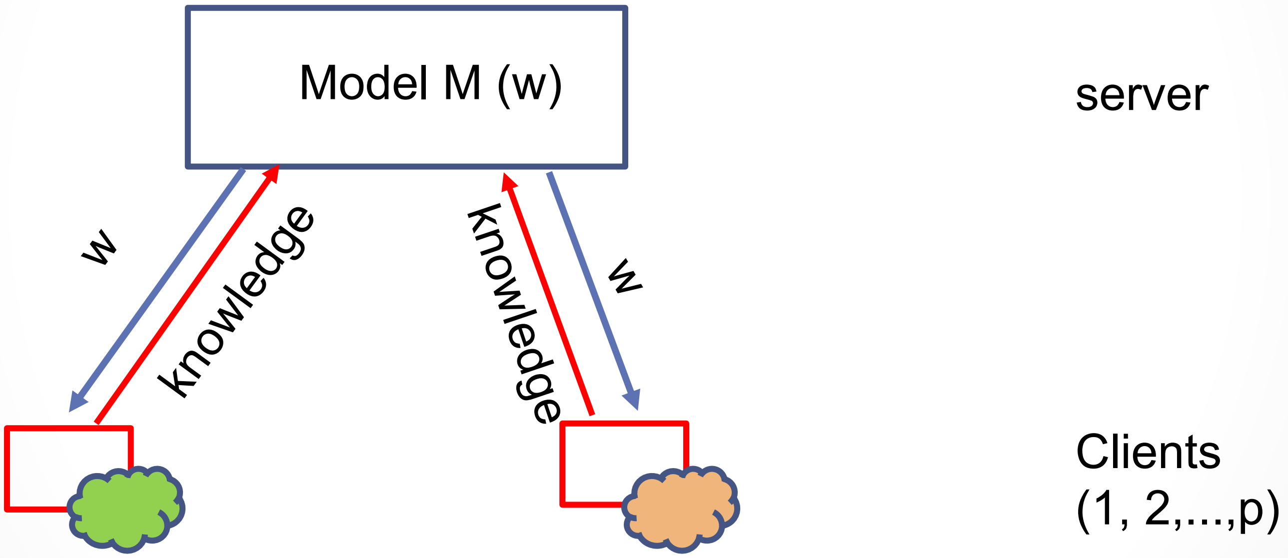

Averaged Federated Learning

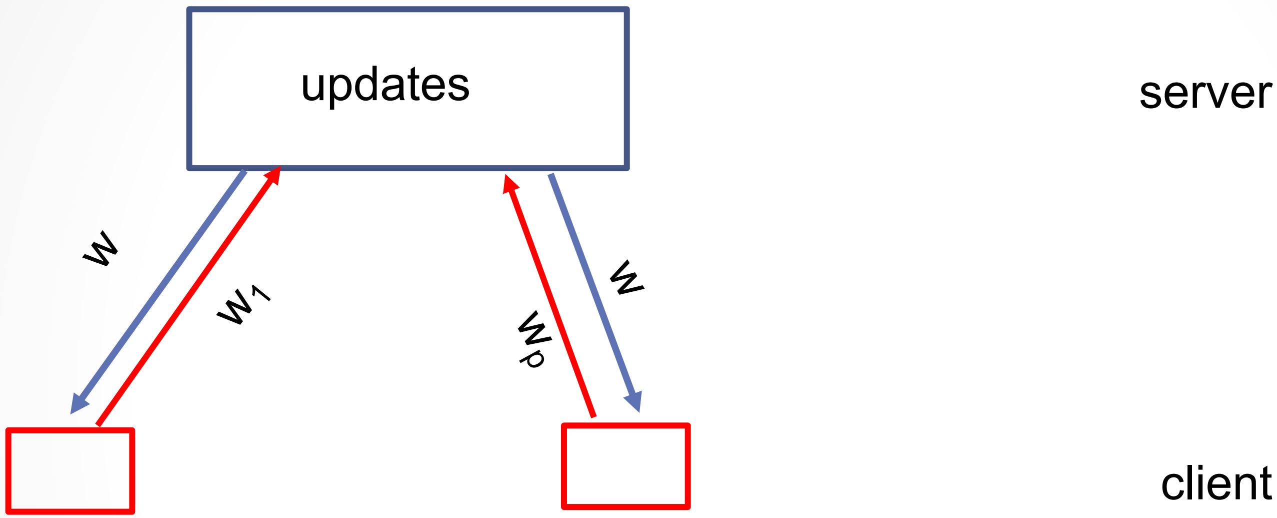

$$\mathrm{w}=\underbrace{\frac{1}{p}\sum_{i=1}^{p}}_{\text{server}}\underbrace{w_i}_{\text{client}}$$

Federated Learning: Gradient-descent learning

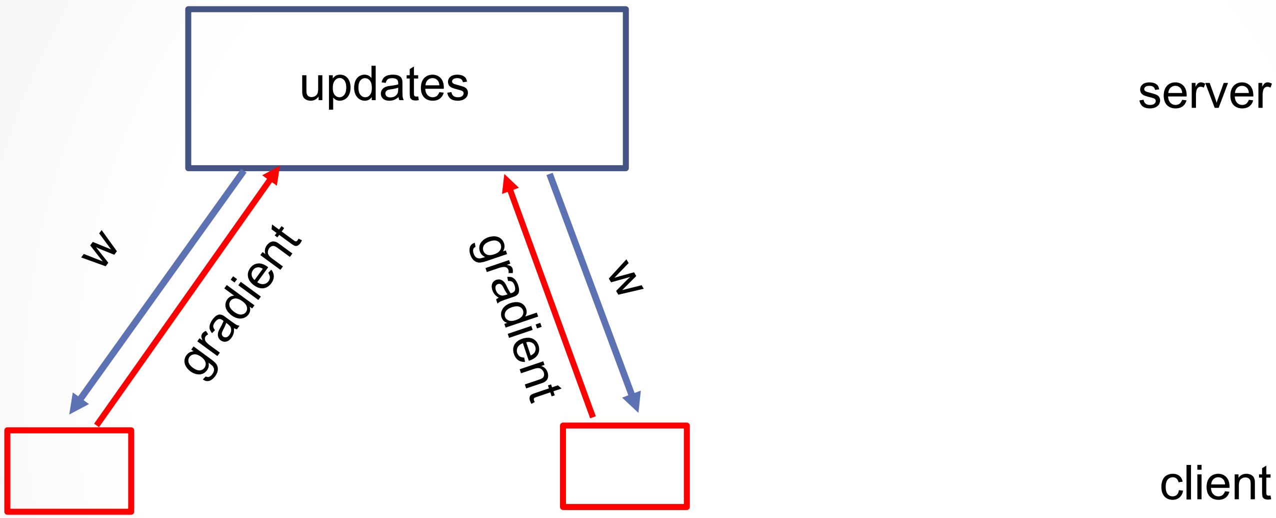

$$\mathbf{w}(iter+1)=\mathbf{w}(iter)\underbrace{-\alpha\sum_{client_i}}_{\text{server}}\underbrace{\text{gradient } Q_i}_{\text{client}}$$

Evaluation of federated learning-based models

Model M confronted with local data D_{ii} of client ii^{th} results in its ***granular*** counterpart $G(M)|_{Dii}$

$G(M)|_{D1}$ $G(M)|_{D2}$ .... $G(M)|_{Dp}$

$G(M)|_{Dii}$ characterized by level of information granularity ε_{ii}

ε_{ii}= arg max(cov*sp)

Granular federated learning-based model- optimization (1)

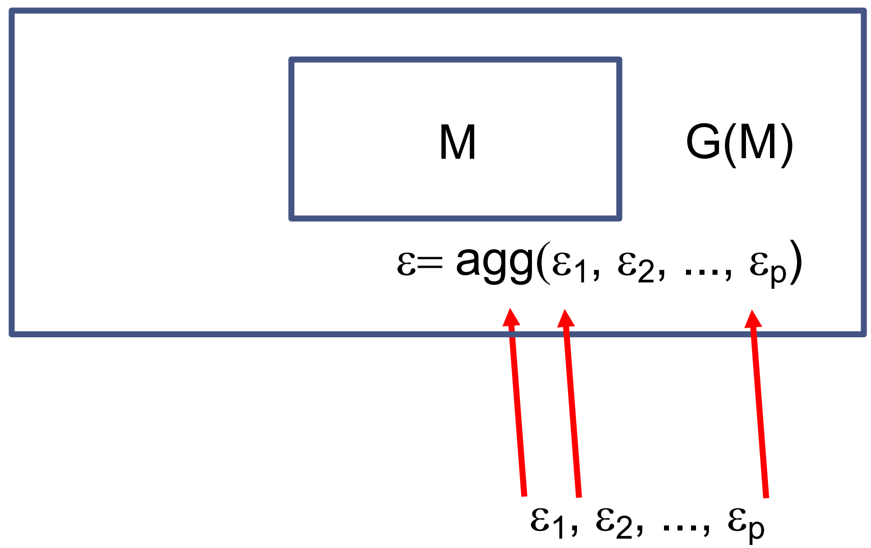

Aggregation of levels of information granularity

$$\varepsilon^* = agg(\varepsilon_1, \varepsilon_2, \ldots, \varepsilon_p, f_1, f_2 \ldots f_p)$$

$f_1, f_2, \ldots f_p$ – weights

agg $\in \mathcal{Agg}$ –family of aggregation operations

Granular federated learning-based model- optimization (2)

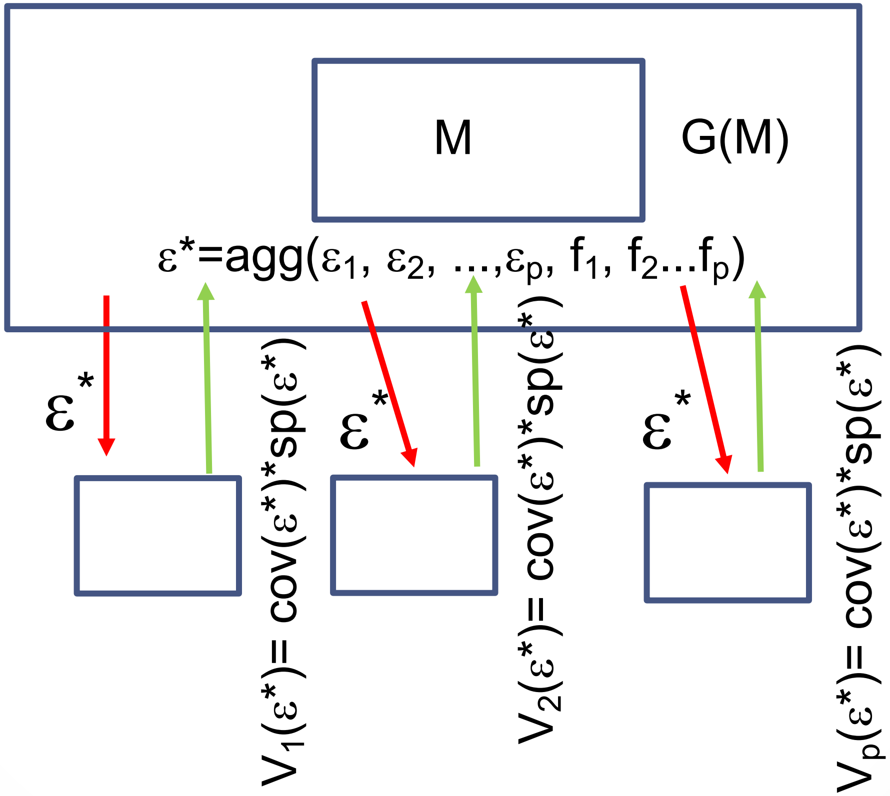

Granular federated learning-based model- optimization (3)

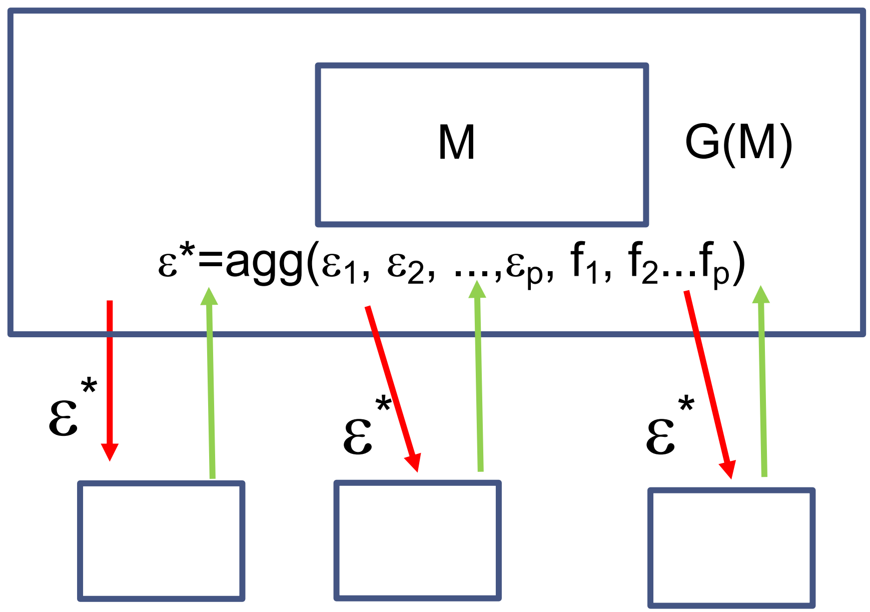

$$(\varepsilon^*_{opt},\ agg_{opt}, f_{1,opt},\ f_{2,opt} \ldots\ f_{p,opt}) =$$

$$= \arg\ \mathrm{Max}_{agg \in \mathcal{A}gg,\ f_1,\ f_2 \ldots f_p} [V_1(\varepsilon *) + V_2(\varepsilon *) + \ldots + V_p(\varepsilon *),\ agg,\ f_1,\ f_2 \ldots f_p)]$$

Aggregation operators: generalized averages

$$agg(a_1, a_2, \ldots, a_n) = \sqrt[p]{\frac{1}{n}\sum_{i=1}^{n}(a_i)^p}$$

$p = 1$ arithmetic mean $agg(a_1, a_2, \ldots, a_n) = \frac{1}{n}\sum_{i=1}^{n}(a_i)$

$p \to 0$ geometric mean $agg(a_1, a_2, \ldots, a_n) = (a_1 a_2 \ldots a_n)^{1/n}$

$p = -1$ harmonic mean $agg(a_1, a_2, \ldots, a_n) = \frac{n}{\sum_{i=1}^{n}(1/a_i)}$

Carbon footprint

Transfer Learning

Credibility of ML models and results

Granular Computing

Transfer learning: an idea

Transfer learning: extraction of previously acquired knowledge and applied to a new similar application

Advantages/motivation:
Small, high quality data
Enhancing robustness of the ML model
Elimination of cold start problem

Terminology
Knowledge reuse, learning by analogy, domain adaptation, Pretraining...

An idea

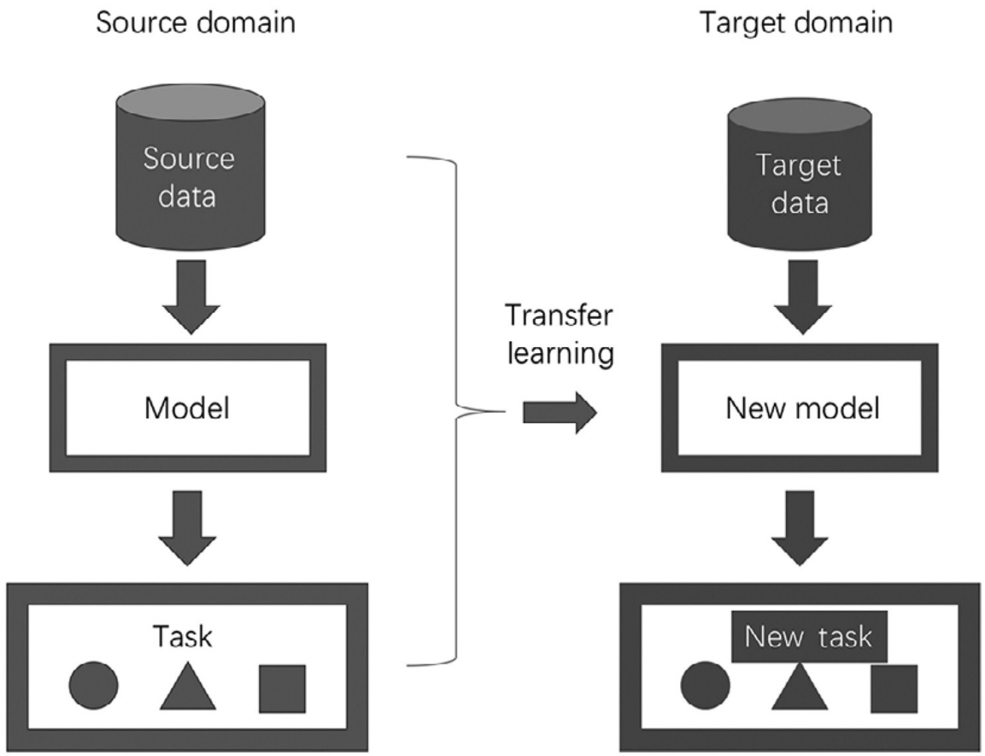

domain: $D_s = \{F_s, P()\}$

task: $T_s=\{Y_s, f_s(.)\}$

domain: $D_t = \{F_t, P()\}$

task: $T_t=\{Y_t, f_t(.)\}$

$D_s \neq D_t$ $T_s \neq T_t$

Transfer Learning with information granules: passive approach

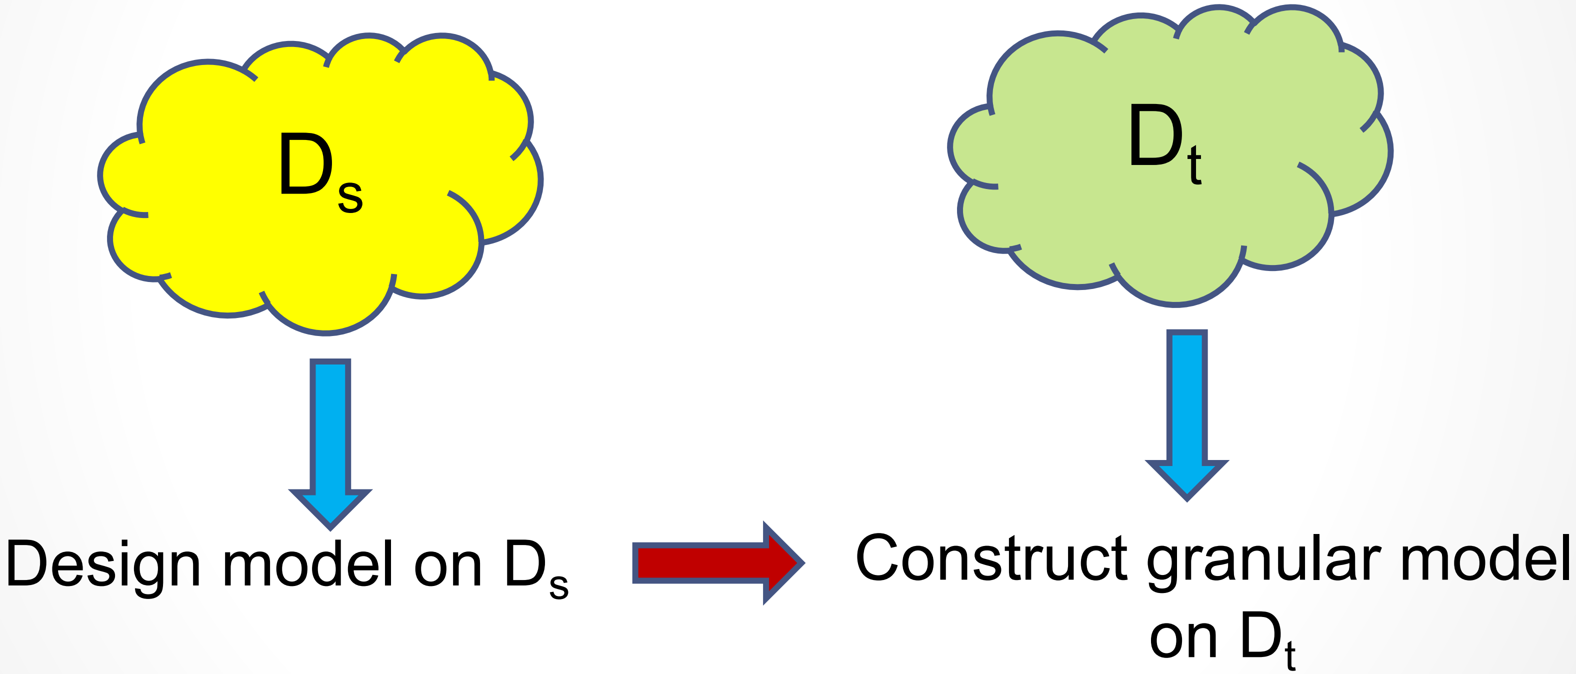

Information granularity associated with model to characterize closeness between source and target domains

Transfer Learning with information granules: active approach

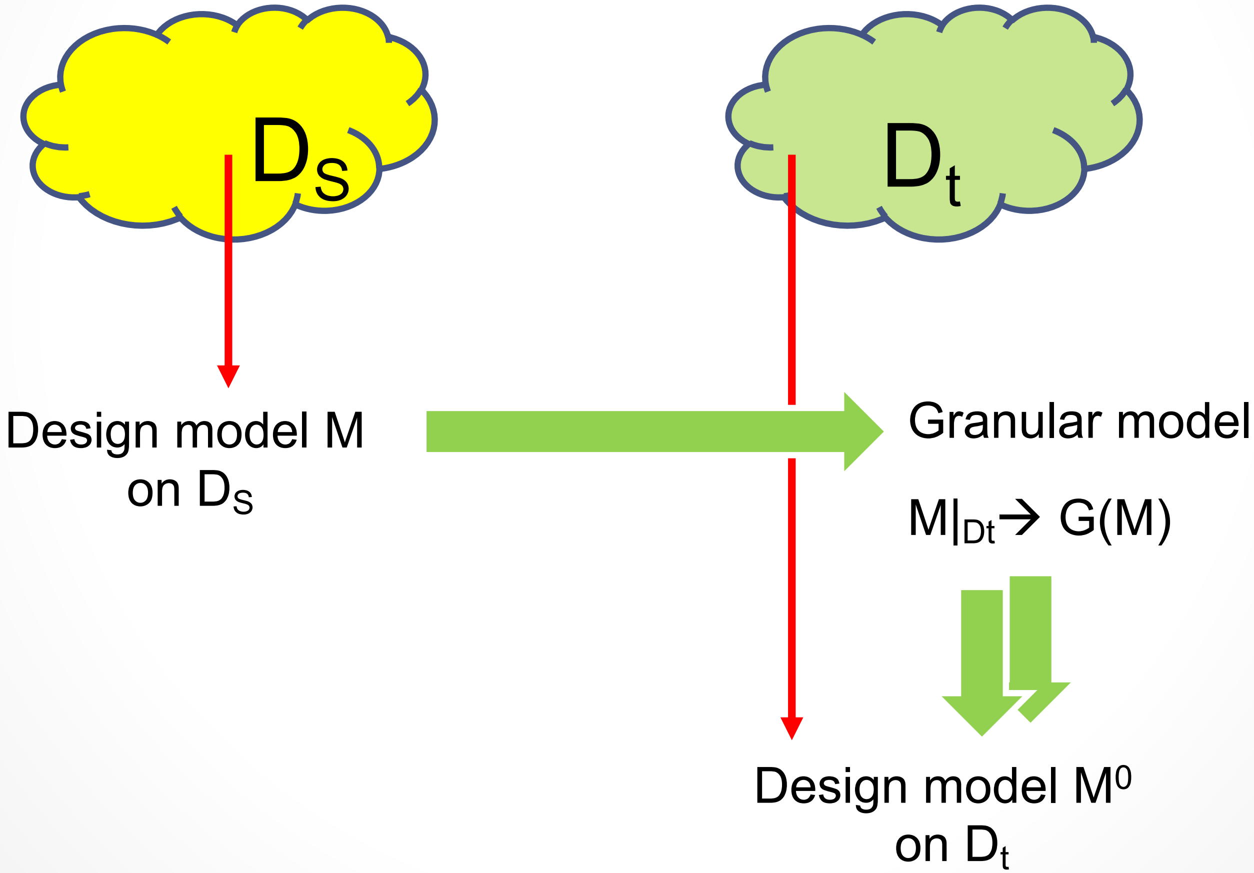

Transfer learning with information granules

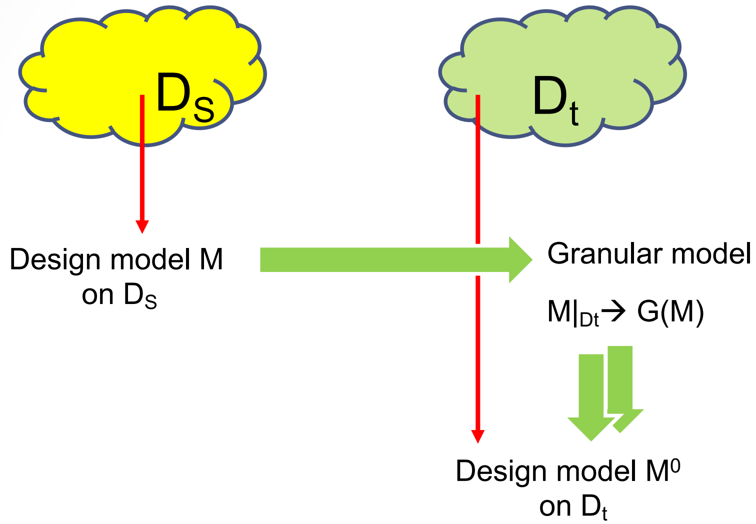

Loss function

$$Q=\sum_{\boldsymbol{D}_t} ||target_k - M^0(\boldsymbol{x}_k, \boldsymbol{w})|| + \alpha \underbrace{\sum_{\boldsymbol{D}_t}[1 - cov(M^0(\boldsymbol{x}_k, \boldsymbol{w}), G(M(\boldsymbol{x}_k))] * sp(G(M(\boldsymbol{x}_k))}_{\textit{granular regularization}}$$

Min $\boldsymbol{w}$ Q $\quad$ $\mathbf{w}=\mathbf{w}-\beta\nabla_{\boldsymbol{w}}Q$

Multisource transfer learning with information granules

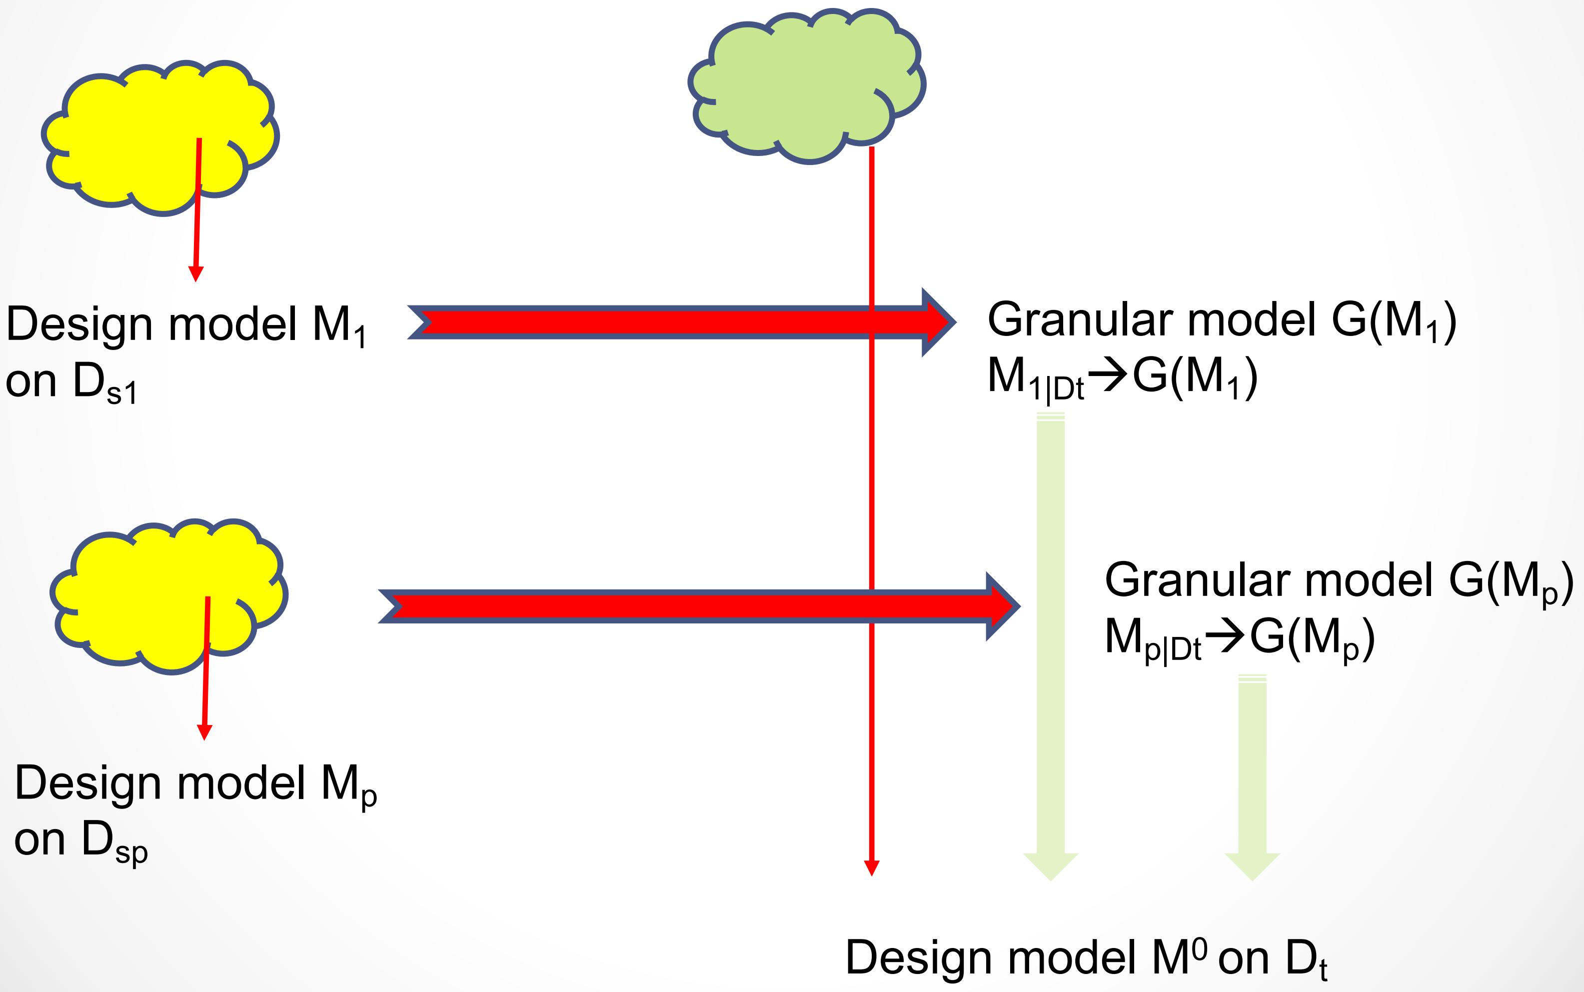

Multisources transfer learning with information granules

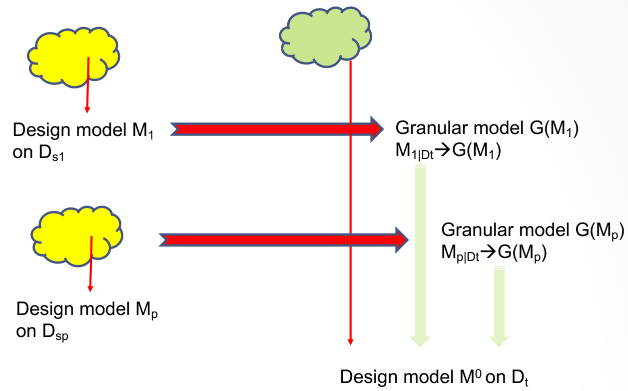

Augmented loss function

$$Q=\sum_{\boldsymbol{D}_t} \| target_k - M^0(\boldsymbol{x}_k, \boldsymbol{w}) \| + \alpha_1 \sum_{\boldsymbol{D}_t} [1 - cov(M^0(\boldsymbol{x}_k, \boldsymbol{w}), G(M_1(\boldsymbol{x}_k))] * sp(G(M_1(\boldsymbol{x}_k)) +$$

$$+\alpha_2 \sum_{D_t} [1 - cov(M^0(\boldsymbol{x}_k, \boldsymbol{w}), G(M_2(\boldsymbol{x}_k))] * sp(G(M_2(\boldsymbol{x}_k)) +$$

$$+\alpha_p \sum_{D_t} [1 - cov(M^0(\boldsymbol{x}_k, \boldsymbol{w}), G(M_p(\boldsymbol{x}_k))] * sp(G(M_p(\boldsymbol{x}_k))$$

granular regularization

Min $\boldsymbol{w}$ Q $\quad$ **w=w**-$\beta \nabla_{\boldsymbol{w}} Q$

Conclusions

New horizons of ML

The role of information granules and Granular Computing

Granular embedding and their role in quantification of results

Future developments: active learning strategies