scch { software competence center hagenberg }

Personalized AI

AI Conference and Expo (AICE) 2024
Sino-Pak Center for Artificial Intelligence (SPCAI)
Pak-Austria Fachhochschule: Institute of Applied Sciences and Technology

Bernhard A. Moser

25.06.2024

AI's Potential, Challenges, and Future Milestones: Insights from Sam Altman

Feb 18,2024

9. Key Milestones in AI Development

    Multimodality: Speech, Images, and Video

    Reasoning Ability

    Reliability

    Customizability and Personalization

    Data Integration and Utilization

Sam Altman
CEO of OpenAI

OpenAI

https://www.toolify.ai/ai-news/ais-potential-challenges-and-future-milestones-insights-from-sam-altman-1414999

General AI - Personalized AI Duality

General AI

domain-independent intelligence
(across cognitive domains)

task versatility

one general model

Foundation Model

Personalized AI

adaptation to individual users and sub-populations

customization to users' needs, preferences, and contexts.

not one-size-fits-all experience

Model Adaptation
Transfer Learning

scch { software competence center hagenberg }

‘Context’ in different contexts

Context in Language

The child picks ____________ in the garden.

Haripur is a _______ in Pakistan.

Students _______ AI at the university of applied sciences.

flowers

study

cooking

tree

town

playing

Approach behind ChatGPT

- Deep Learning Modell mit „attention“ Mechanismus (2014)
- **g**enerative **p**re-trained **t**ransformers (2017)

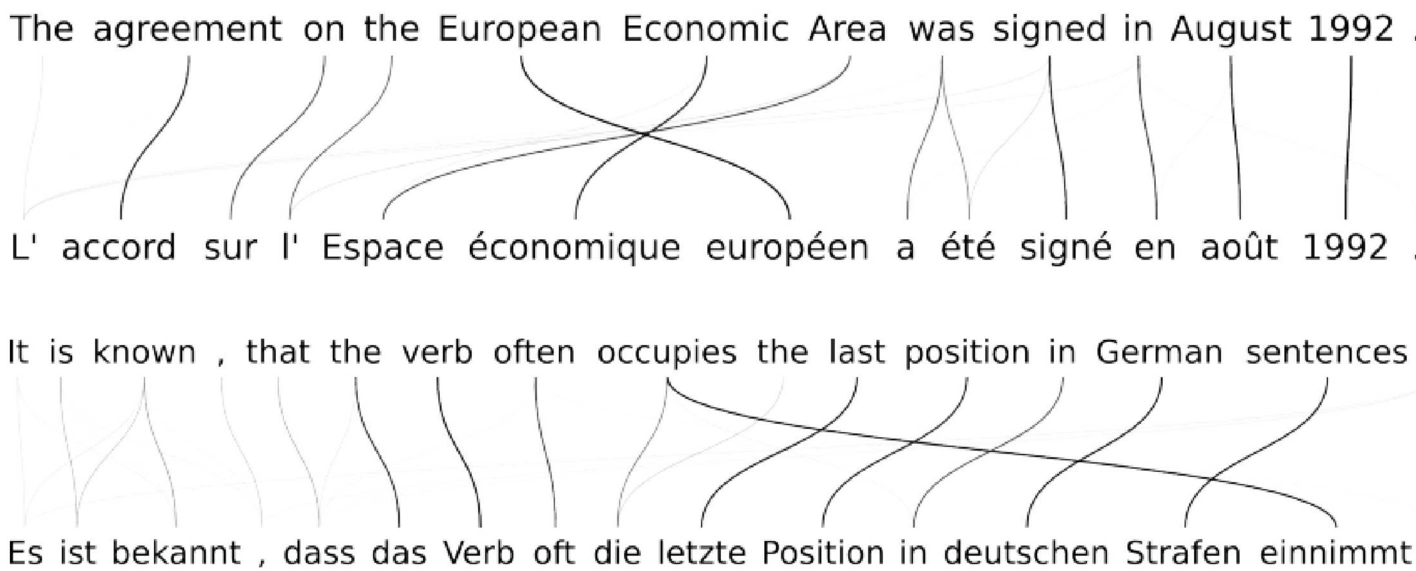

[penalty???]

D. Bahdanu, K Cho, Y. Bengio: Neural Machine Translation by Jointly Learning to Align and Translate, ICRL 2015.

Development since then

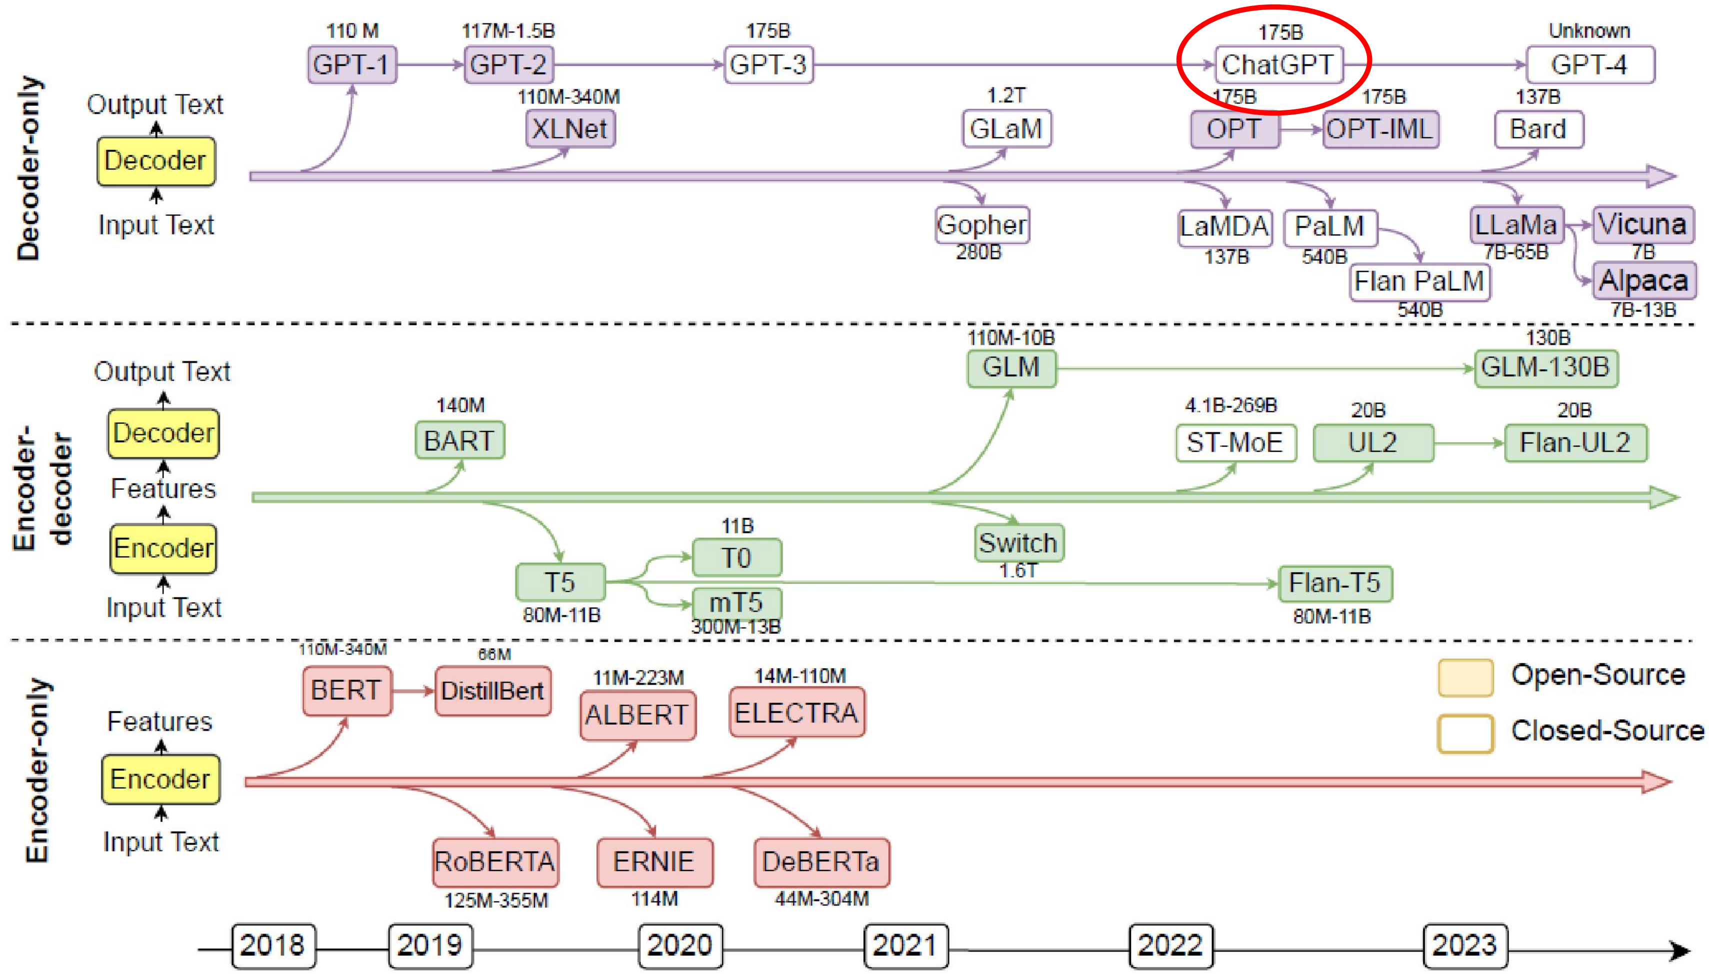

Humans in Process-Centric Scenarios

This project receives funding in the European Commission's Horizon 2020 Research Programme under Grant Agreement Number 957402

https://www.teamingai-project.eu/

TEAMING.AI hybrid AI approach

- create knowledge and process model
  - digital shadow of dynamic enviroment
- context tracking
  - to supplement unobserved parameters with knowledg
  - to ensure transparency and exaplainability
  - to increase situational awareness
  - to check compliance with regulations
- ML on dynamic knowledge graph
  - to enable self-learning and adaptability
- Human/team on/in the loop
  - to guarantee human oversight
  - to resolve ambiguities in situational awareness and decision making
  - to accelerate phases of transition (setup, maintenance)

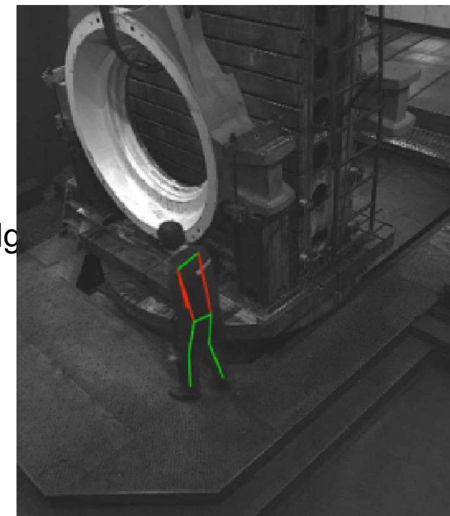

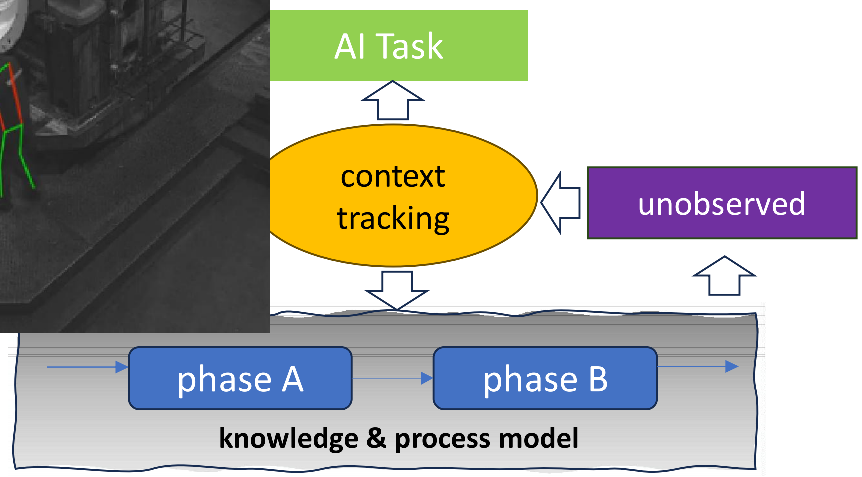

Context by Biased Sampling

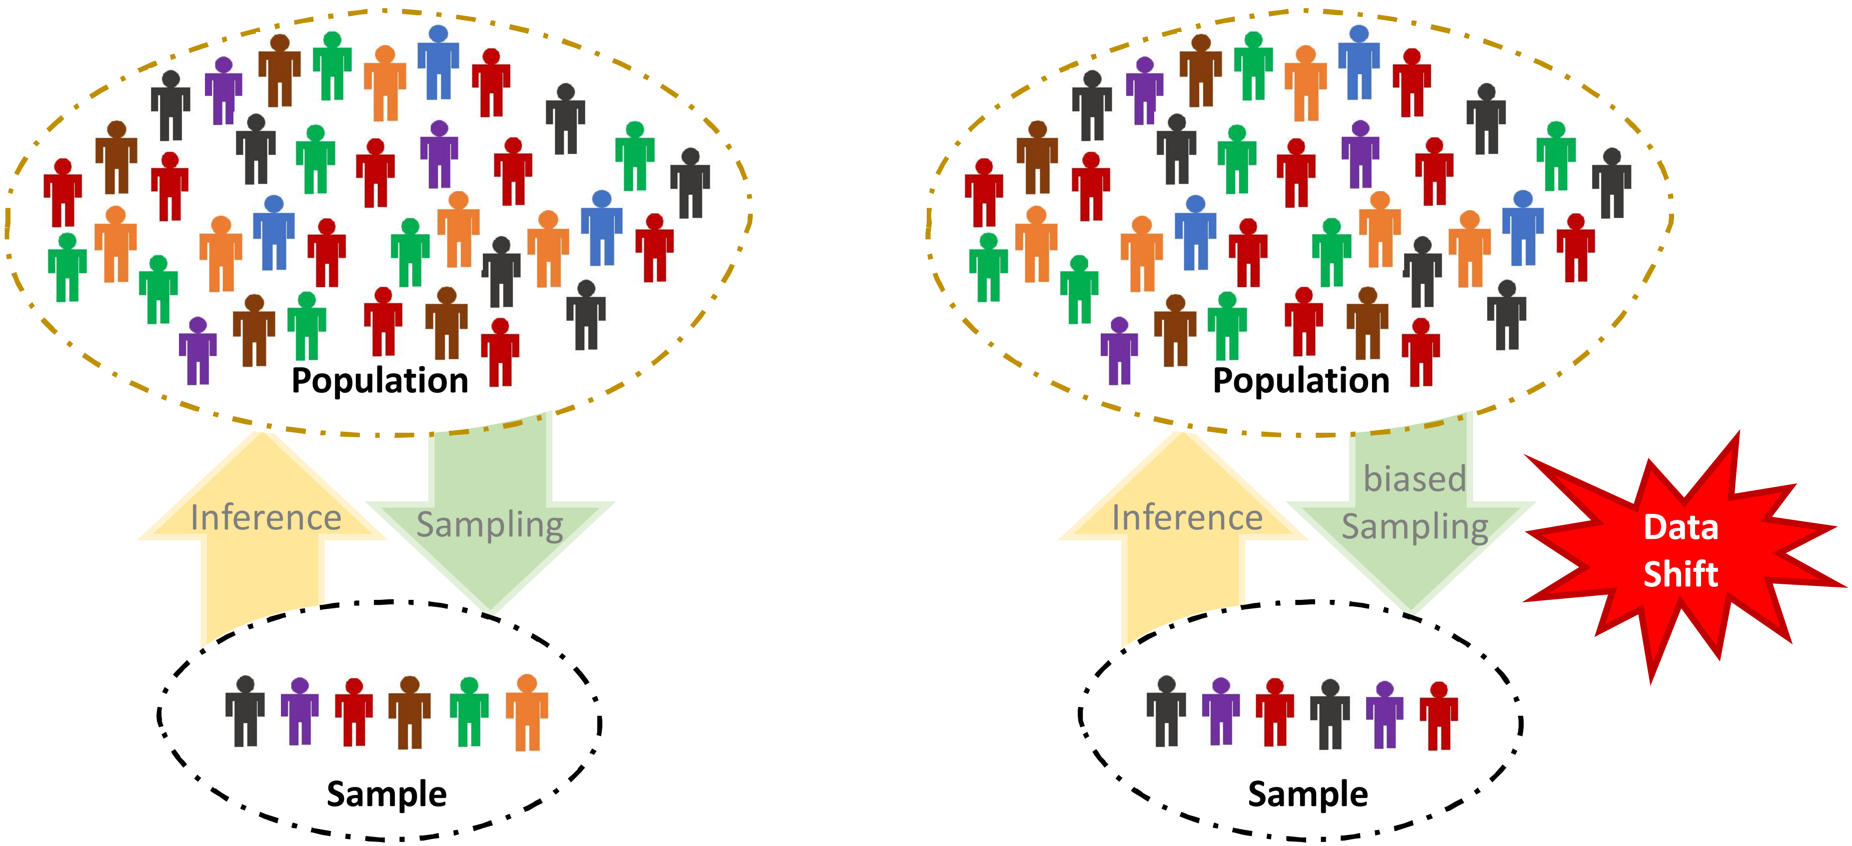

Breakthrough with Deep Learning

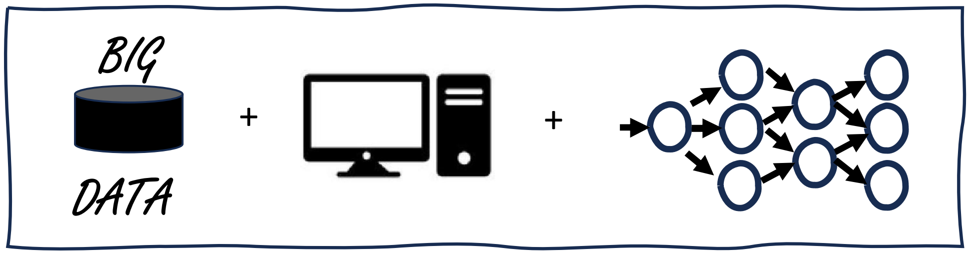

- 2006: „Microsoft Research“
- 2011: „Google Brain“
- 2012: breakthrough in visual recognition
  - Backpropagation-Algorithm (G. Hinton)
- 2013: „Facebook AI Research“
- 2016: LSTMs (S. Hocheiter & J. Schmidhuber/1997): breakthrough in NLP tasks
- 2016: Google‘s ALPHA GO

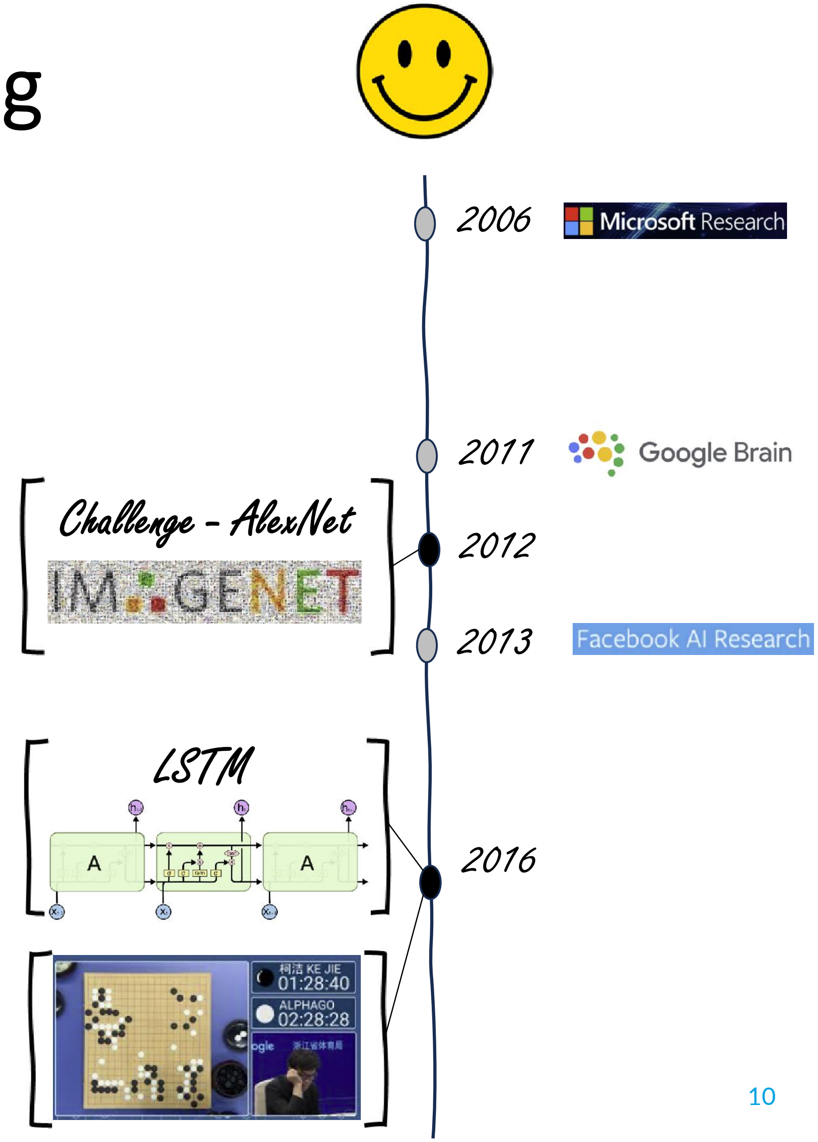

From ImageNet to ObjectNet

ImageNet

- > 14 Mio annotated images for training
- 50k images for testing
- benchmark in image classification and object detection
- ImageNet Large Scale Visual Recognition Challenge (ILSVRC)

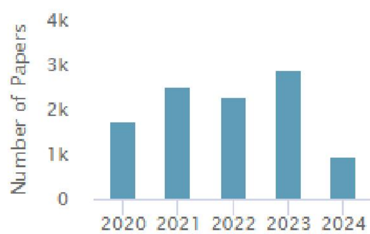

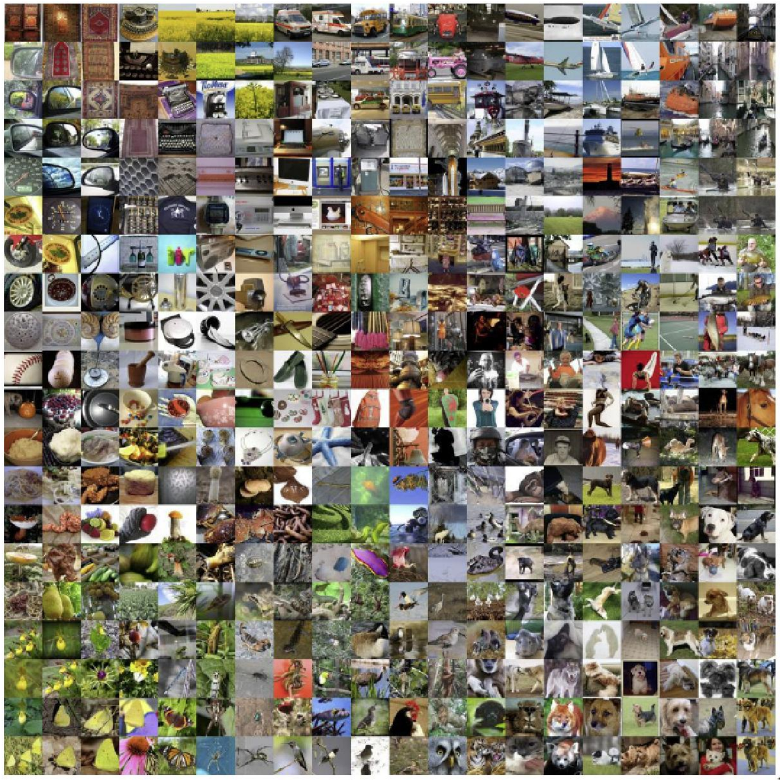

https://paperswithcode.com/dataset/imagenet

From ImageNet to ObjectNet

ObjectNet

- no training set, only a test set!
- collected to intentionally show objects from new viewpoints on new backgrounds
- with controls for rotation, background, and viewpoint
- ObjectNet is the same size as the ImageNet test set (50k images)
- 313 object classes with 113 overlapping ImageNet
- The dataset is both easier than ImageNet – objects are largely centered and unoccluded – and harder, due to the controls

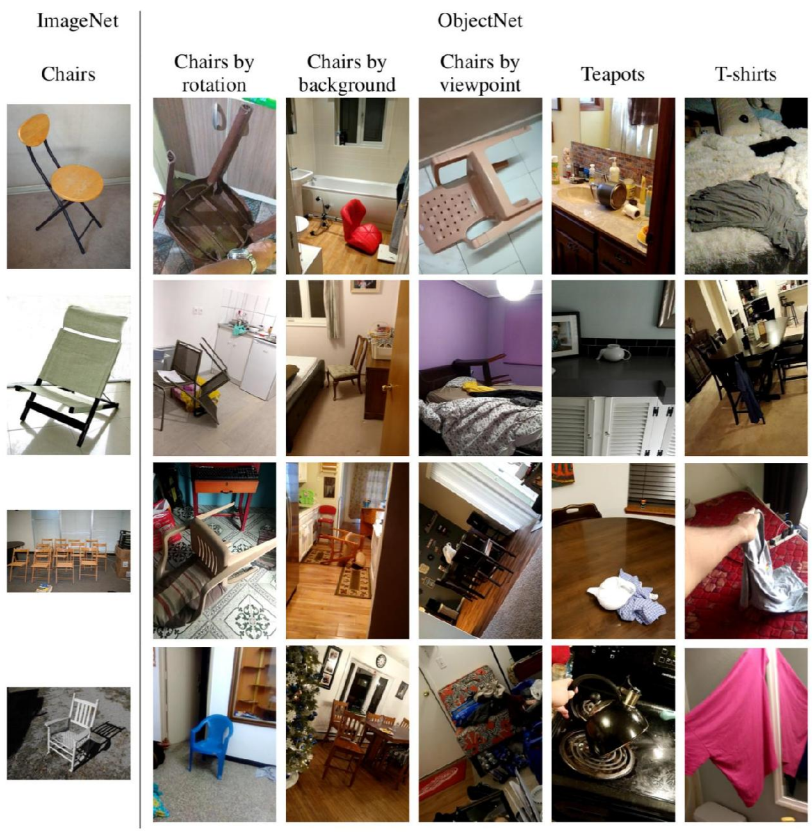

https://objectnet.dev/

Easy for Humans, Hard for Machines

- Large performance drop, what you can expect from vision systems in the real world!
- predictive of real-world performance

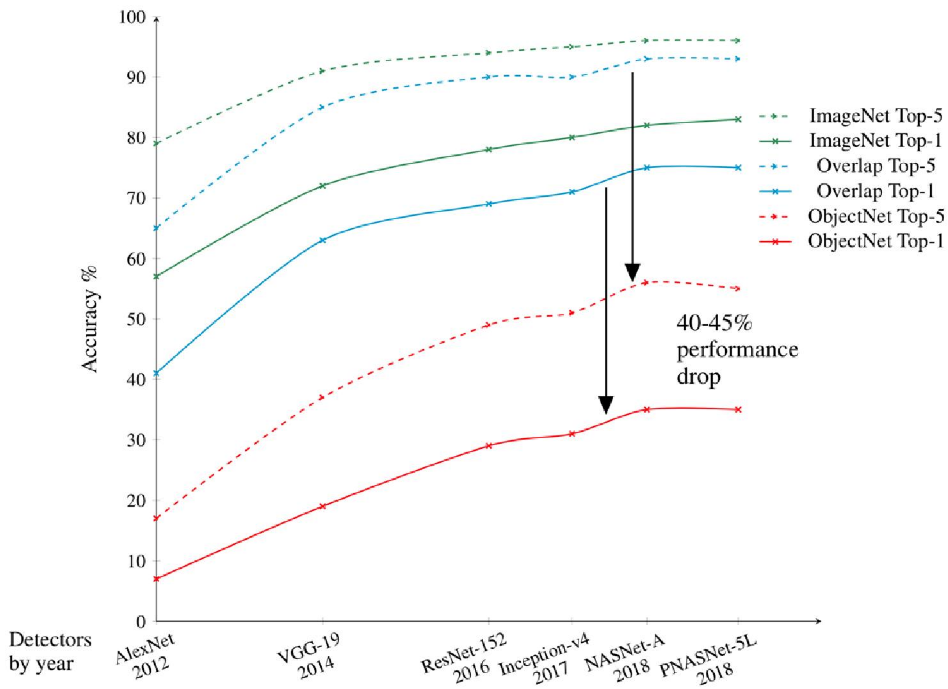

Big Data?

- ImageNet is a biased sample from real world images
- Actually, in real world applications we only have limited data!
- This a fundamental issue
  - for a learning theory
  - classical statistical learning theeory needs to be extended

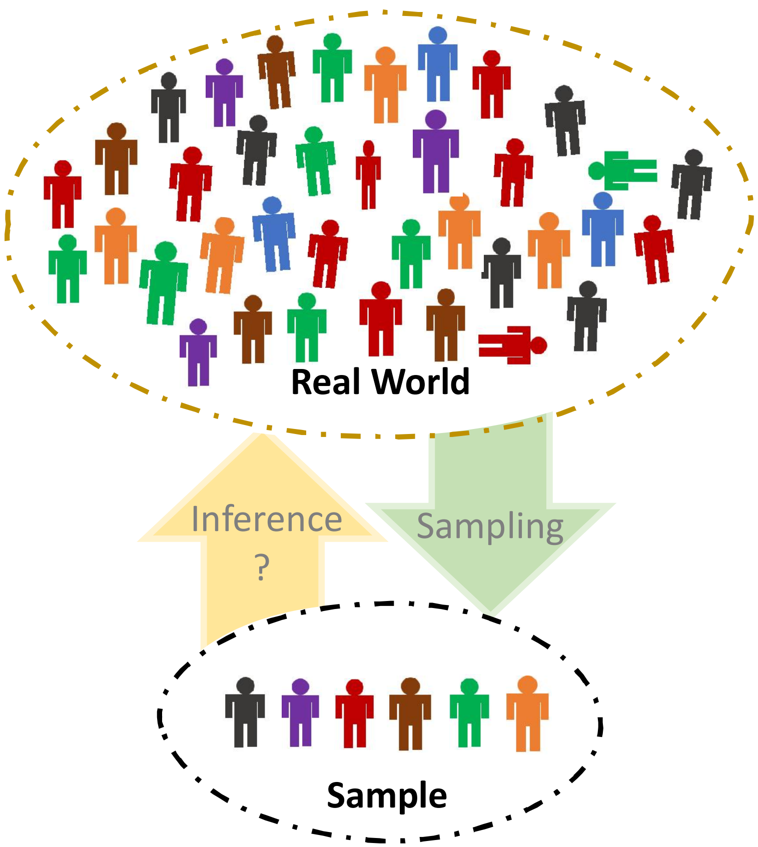

sampling is not representative

Application Scenarios for Costomized AI

- Spam filter
  - different users have different email statistics

- Personalized Medicine
  - Diabetes prediction (changed stress level)
  - Tumor cell segmentation (changed gene expression)

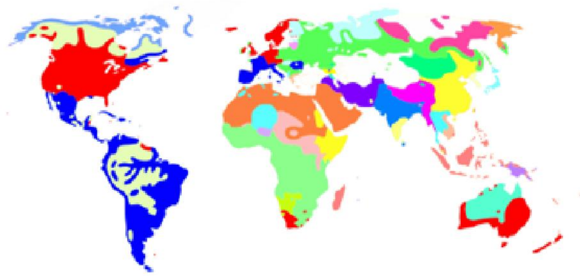

Gene Diversity Map

- Industry
  - Quality control (changed product features or production lines)
  - Robotics (special views on objects)
  - Occupational safety and health (changing processes)
  - New business models

New Business Models

https://www.s3ai.at/

Requirements from Industry

- business obstacles
    - lack of effective business models
    - security concerns
- data/ML problems
    - lack of big data, shift in distribution

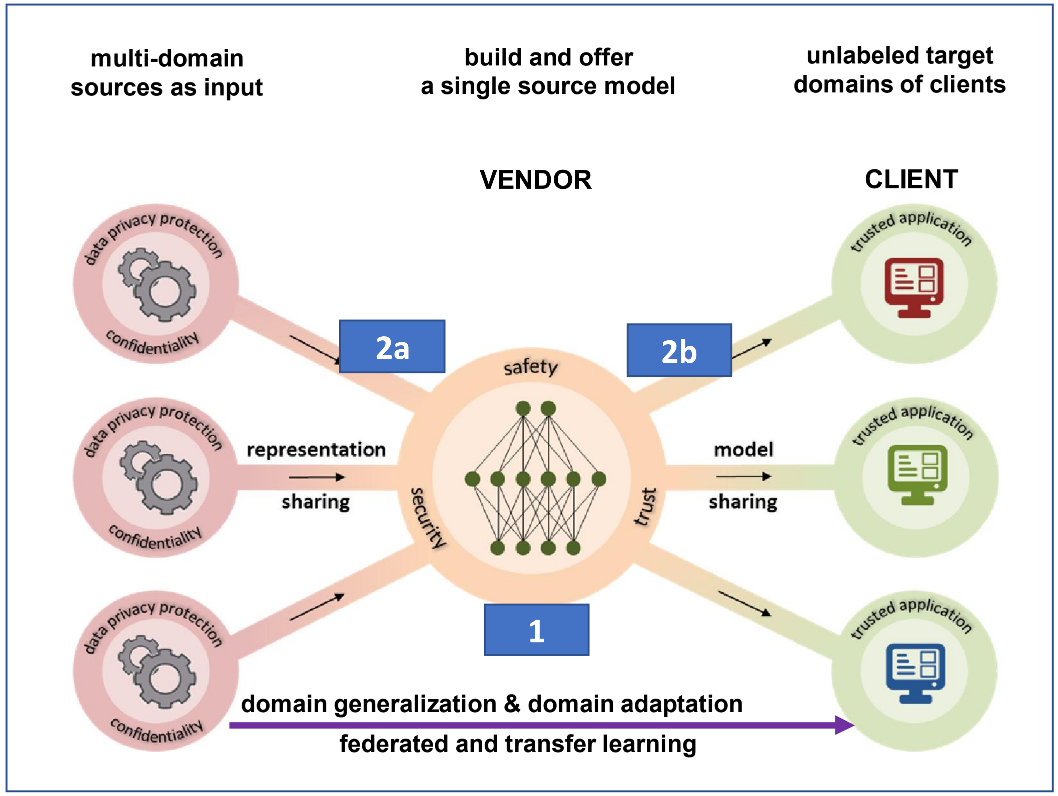

vendor-client shared AI business model

scch {
software
competence
center
hagenberg
}

Tackling Data Shift

Domain Shift

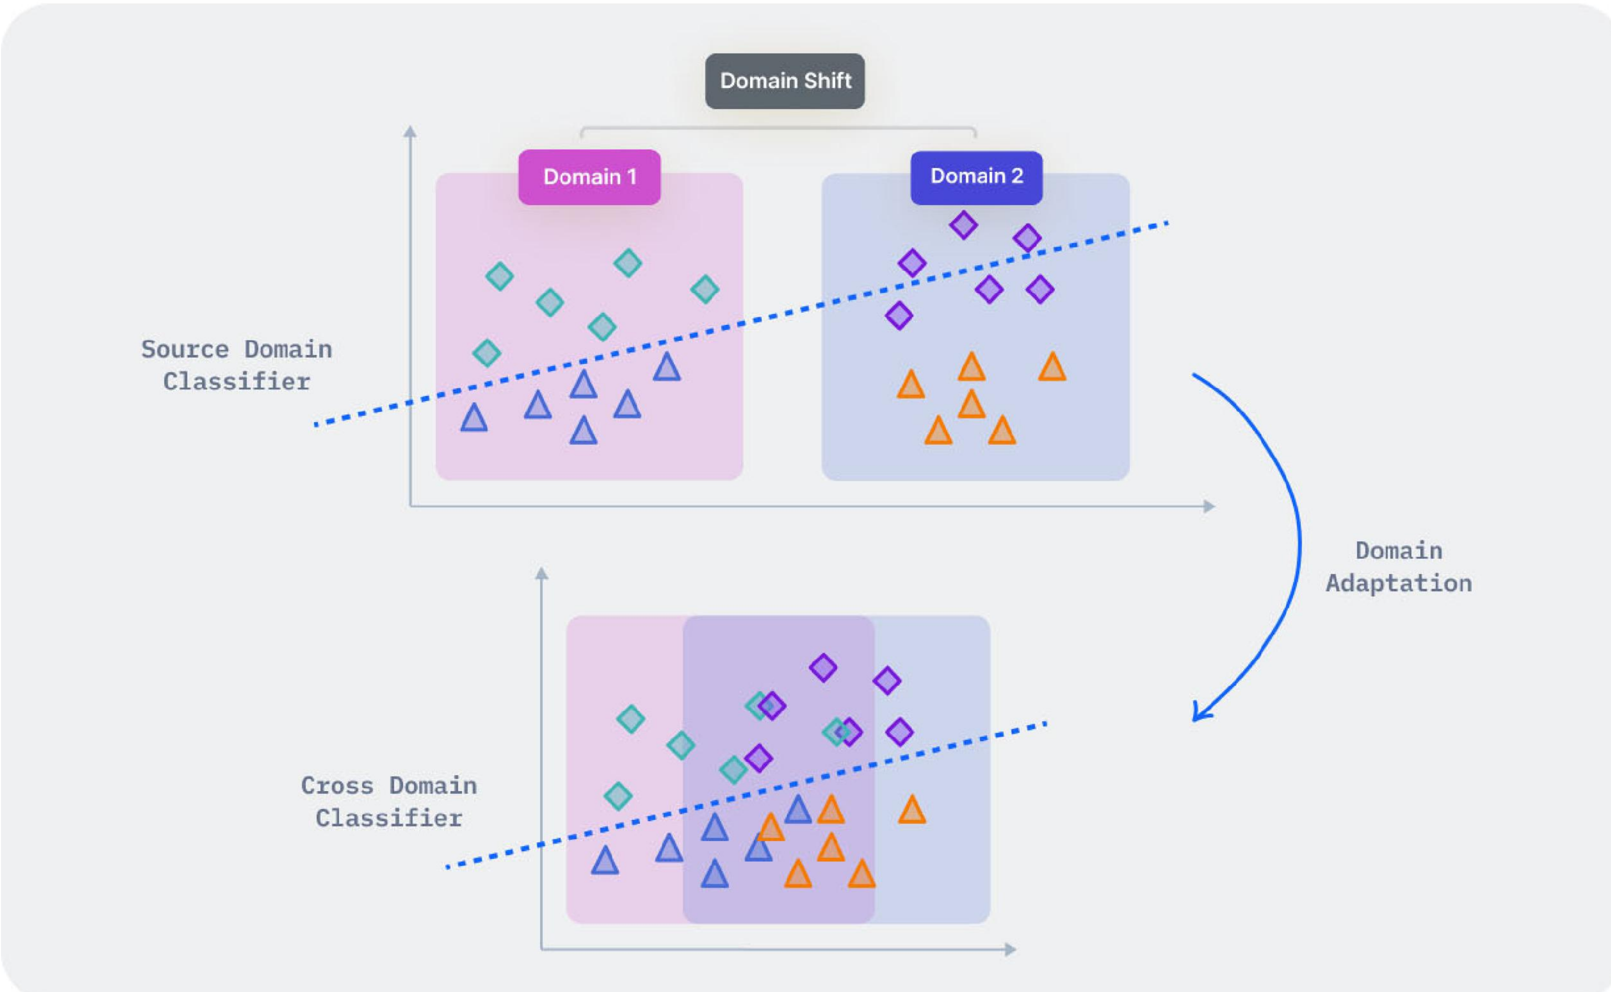

Limitations / Challenges

- The success of the transfer heavily depends on the similarity between the source and target tasks.
- If the tasks are too dissimilar, the transferred knowledge may not be beneficial and could even harm the performance on the target task.
- Fine-tuning a pre-trained model requires careful hyperparameter tuning to avoid catastrophic forgetting of the knowledge learned from the source task.

ICLR

Marius-Constantin Dinu, Markus Holzleitner, Maximilian Beck, Hoan Duc Nguyen, Andrea Huber, Hamid Eghbal-zadeh, Bernhard A. Moser, Sergei Pereverzyev, Sepp Hochreiter, Werner Zellinger

Addressing Parameter Choice Issues in Unsupervised Domain Adaptation by Aggregation

ICLR2023 (notable top 5%)

Domain Adaptation by Aggregation

Source data

$$\{(x_i, y_i)\}_{i=1}^{s} \sim p$$

Target data w/o labels

$$\{x_i'\}_{i=1}^{t} \sim q_X$$

[Peng et al. 2019]

Goal: Learn model $f : X \to Y \subset \mathbb{R}^d$ with small error

$$\mathcal{E}_q(f) := \int_{X \times Y} \|f(x) - y\|_Y^2 \, \mathrm{d}q(x, y)$$

Problem

$$A_{\boxed{\alpha,\lambda,\rho,\ldots}} : \left(\{x_i, y_i\}_{i=1}^{s}, \{x'_i\}_{i=1}^{t}\right) \mapsto (f : X \to Y)$$

- α learning rate
- λ loss weights
- $\ldots$

How to choose parameters w/o target labels?

State of the Art - Step 1

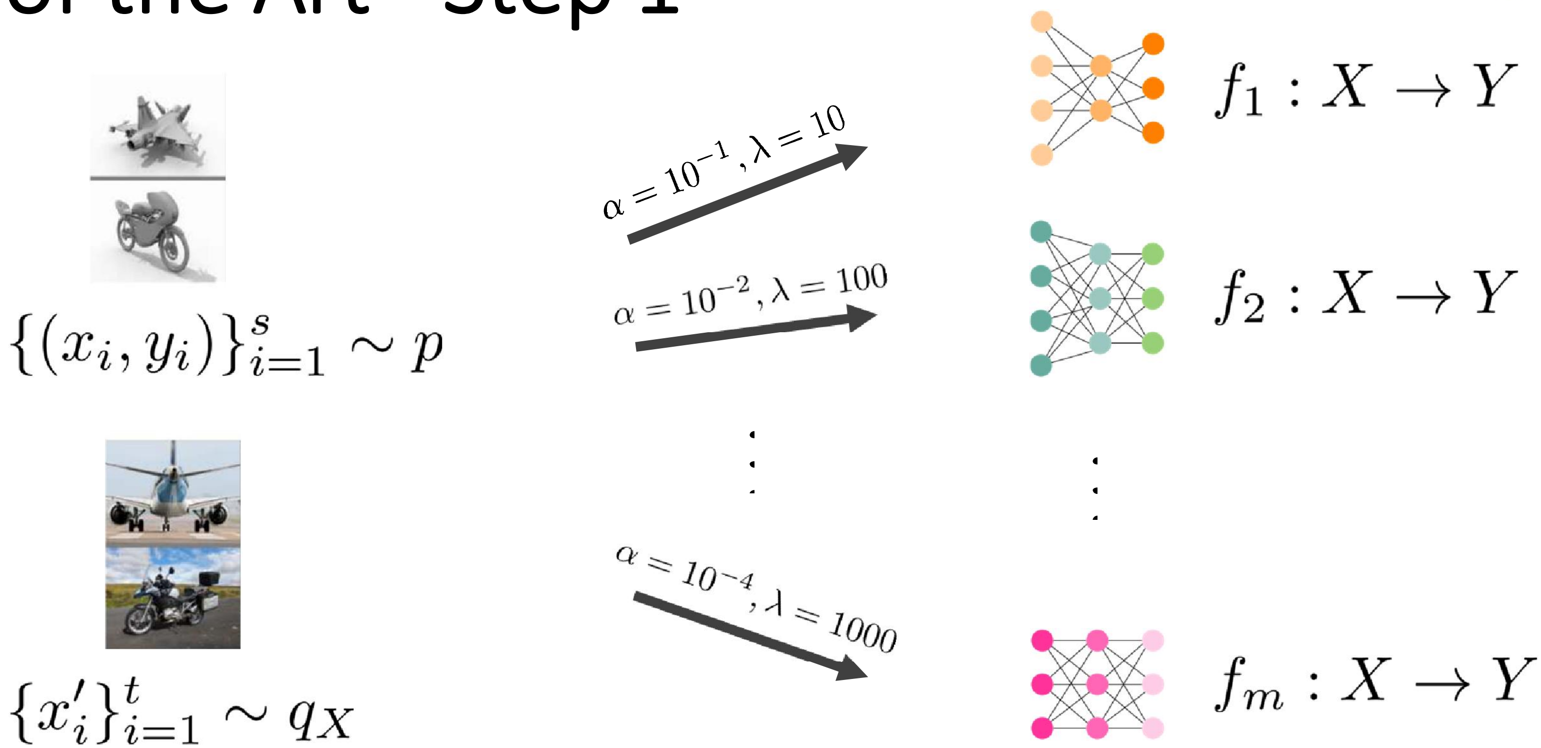

Compute different models $f_1, \ldots, f_m : X \to Y$ by running algorithm with different parametrizations.

State of the Art - Step 2

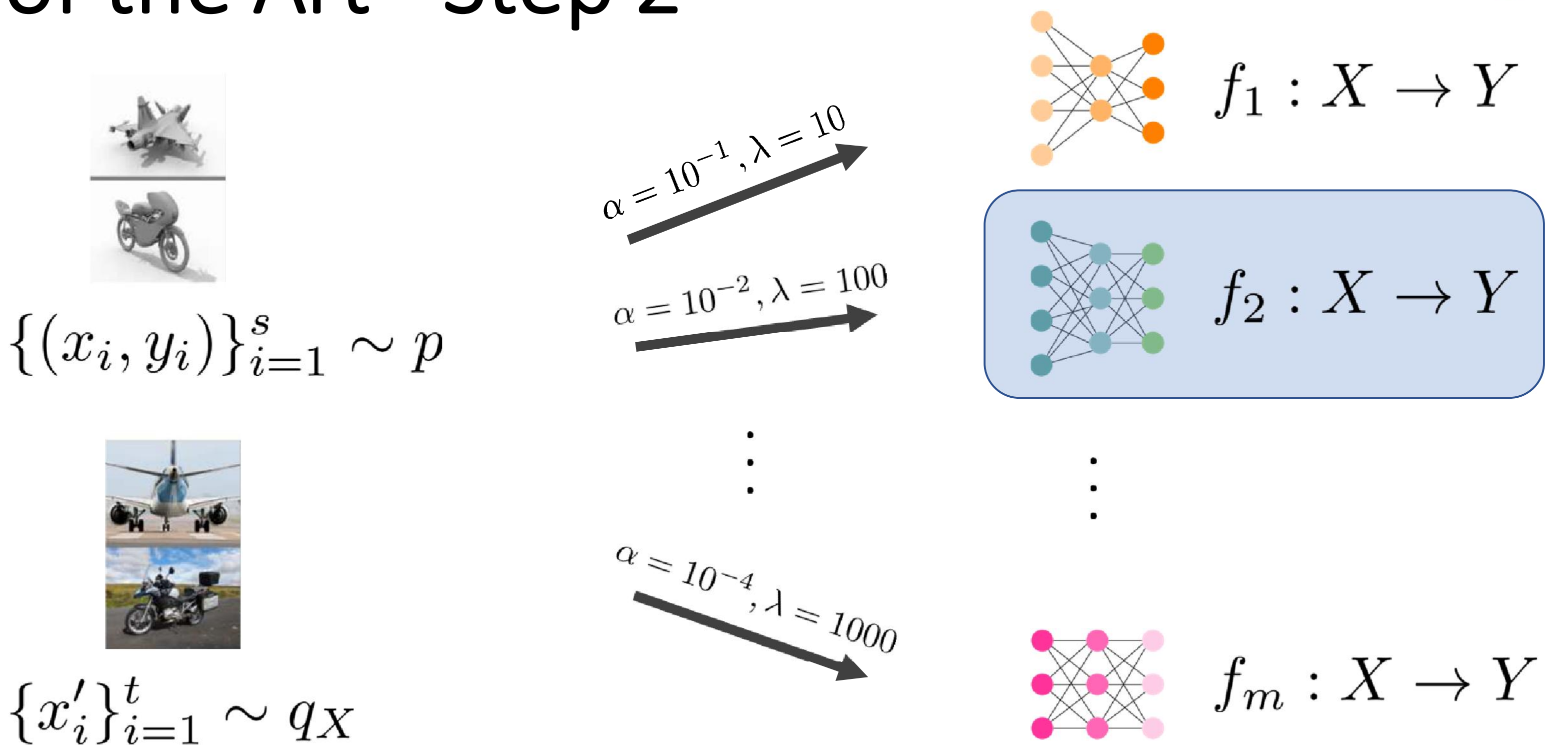

Select model $f^{\text{sel}} := \underset{f \in \{f_1, \ldots, f_m\}}{\arg\min} \; \mathcal{E}_q(f)$ with smallest error.

Our Approach

Compute linear aggregation $f^{\text{agg}} := \sum_{i=1}^{m} c_i f_i$ with

$$\mathcal{E}_q\left(f^{\text{agg}}\right) = \min_{c_1,\ldots,c_m \in \mathbb{R}} \mathcal{E}_q\Big(\sum_{i=1}^{m} c_i f_i\Big)$$

Error is smaller than best single model

$$\mathcal{E}_q(f^{\text{agg}}) \leq \mathcal{E}_q(f^{\text{sel}})$$

Optimization: Vector-Valued Least Squares

perfect solution

$$c^{\mathrm{agg}} := G^{-1} g = \underset{(c_1,\ldots,c_m)\in\mathbb{R}^m}{\arg\min} \int_X \left\| \sum_{i=1}^m c_i f_i(x) - \boxed{f_q(x)} \right\|_Y^2 \mathrm{d}q_X(x)$$

with Bayes predictor and Gram matrix

$$f_q(x) = \int_Y y \, \mathrm{d}q(y|x) \qquad G = \left(\int_X \langle f_k(x), f_u(x) \rangle_Y \, \mathrm{d}q_X(x) \right)_{k,u=1}^m$$

and vector

$$g = \left(\int_X \langle f_q(x), f_k(x) \rangle_Y \, \mathrm{d}q_X(x) \right)_{k=1}^m$$

Optimization (contd)

$$c^{\text{agg}} := G^{-1}g = \underset{(c_1,\dots,c_m)\in\mathbb{R}^m}{\arg\min} \int_X \left\| \sum_{i=1}^{m} c_i f_i(x) - \boxed{f_q(x)} \right\|_Y^2 \mathrm{d}q_X(x)$$

with Bayes predictor and Gram matrix

$$\boxed{f_q(x) = \int_Y y \,\mathrm{d}q(y|x)} \qquad G = \left(\int_X \langle f_k(x), f_u(x) \rangle_Y \,\mathrm{d}q_X(x) \right)_{k,u=1}^{m}$$

and vector

not estimable!

$$g = \left(\int_X \langle \boxed{f_q(x)}, f_k(x) \rangle_Y \,\mathrm{d}q_X(x) \right)_{k=1}^{m}$$

Solution: Importance Weighting

Under assumptions

- covariate shift $p(y|x) = \boxed{q(y|x)}$
- bounded density ratio $\beta(x) := \frac{\mathrm{d}q_X}{\mathrm{d}p_X}(x) \in [0, B]$

we get

$$g = \left(\int_X \langle f_q(x), f_k(x) \rangle_Y \, \mathrm{d}q_X(x) \right)_{k=1}^m$$

$\mathrm{d}q_X$

$$f_q(x) = \int_Y y \, \boxed{\mathrm{d}q(y|x)}$$

$$g = \left(\int_X \langle f_p(x), f_k(x) \rangle_Y \, \beta(x) \, \mathrm{d}p_X(x) \right)_{k=1}^m$$

[Shimodaira 2000, Kanamori et al. 2009]

Solution: Importance Weighting

Under assumptions

- covariate shift $p(y|x) = \boxed{q(y|x)}$
- bounded density ratio $\beta(x) := \frac{\mathrm{d}q_X}{\mathrm{d}p_X}(x) \in [0, B]$

we get

$$g = \left(\int_X \langle f_q(x), f_k(x) \rangle_Y \boxed{\mathrm{d}q_X(x)} \right)_{k=1}^{m}$$

$$f_q(x) = \int_Y y \boxed{\mathrm{d}q(y|x)}$$

$$\tilde{g} \longleftarrow \quad g = \left(\int_X \langle \boxed{f_p(x)}, f_k(x) \rangle_Y \boxed{\beta(x)\, \mathrm{d}p_X}(x) \right)_{k=1}^{m}$$

[Shimodaira 2000, Kanamori et al. 2009]

New Algorithm

Step 1: Estimate density ratio $\beta(x)$, e.g. by [Sugiyama et al. 2012]

Step 2: Compute aggregation $\widetilde{f} = \sum_{i=1}^{m} \widetilde{c}_i f_i$ with $\widetilde{c} := \widetilde{G}^{-1}\widetilde{g}$

$$\widetilde{g} = \left(\frac{1}{s} \sum_{i=1}^{s} \beta(x_i) \langle y_i, f_k(x_i) \rangle_Y \right)_{k=1}^{m} \qquad \widetilde{G} = \left(\frac{1}{t} \sum_{i=1}^{t} \langle f_k(x_i'), f_u(x_i') \rangle_Y \right)_{k,u=1}^{m}$$

Result 1: Convergence Rate

With probability at least $1-\delta$ for large enough s and t,

$$\left\|\widetilde{f}-f_q\right\|_{L^2(q_X)} \leq \min_{c_1,\ldots,c_m\in\mathbb{R}} \left\|\sum_{i=1}^{m} c_i f_i - f_q\right\|_{L^2(q_X)} + C\left(s^{-\frac{1}{2}}+t^{-\frac{1}{2}}\right)\log^{\frac{1}{2}}\left(\frac{1}{\delta}\right)$$

Sample size of source, resp. target

Empirical Performance

Methods
HoMM
AdvSKM
DIRT
DDC
CMD
MMDA
CoDATS
Deep-Coral
CDAN
DANN
DSAN

Amazon Reviews

Text

DomainNet

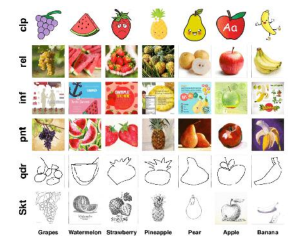

Images

HHAR

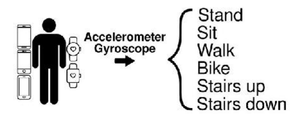

Sleep-EDF

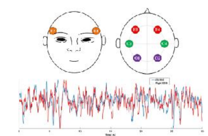

UCI-HAR

WISDM

Time-Series

		Heuristic				Theoretical error guarantees			
Dataset	SO	TMV	TMR	TCR	SOR	IWV	DEV	IWA (ours)	TB
Transformed Moons	0.989(±0.008)	0.980(±0.006)	0.981(±0.007)	*0.997(±0.002)*	0.989(±0.010)	0.989(±0.008)	0.981(±0.022)	**0.997(±0.002)**	0.997(±0.005)
Amazon Reviews	0.767(±0.011)	0.787(±0.009)	0.786(±0.010)	0.786(±0.010)	*0.789(±0.010)*	0.772(±0.014)	0.764(±0.019)	**0.788(±0.009)**	0.781(±0.012)
MiniDomainNet	0.507(±0.022)	*0.526(±0.011)*	0.525(±0.014)	0.526(±0.013)	0.518(±0.012)	0.513(±0.022)	0.515(±0.028)	**0.531(±0.011)**	0.534(±0.022)
Sleep-EDF	0.655(±0.054)	*0.729(±0.018)*	0.729(±0.024)	0.725(±0.023)	0.717(±0.028)	0.700(±0.052)	0.660(±0.057)	**0.737(±0.020)**	0.712(±0.045)
UCI-HAR	0.770(±0.046)	*0.840(±0.017)*	0.833(±0.023)	0.832(±0.024)	0.769(±0.060)	0.774(±0.070)	0.765(±0.090)	**0.835(±0.020)**	0.850(±0.029)
HHAR	0.732(±0.042)	0.771(±0.015)	0.768(±0.017)	*0.771(±0.018)*	0.722(±0.068)	0.746(±0.037)	0.722(±0.063)	**0.787(±0.012)**	0.784(±0.028)
WISDM	0.736(±0.050)	*0.768(±0.027)*	0.768(±0.036)	0.765(±0.037)	0.737(±0.062)	0.736(±0.052)	0.726(±0.077)	**0.764(±0.025)**	0.771(±0.046)

Ragab et al., ADATIME: A Benchmarking Suite for Domain Adaptation on Time Series Data. ACM Trans. Knowl. Discov. Data 2023, https://doi.org/10.1145/3587937

Open question: When to adapt?

When to adapt?

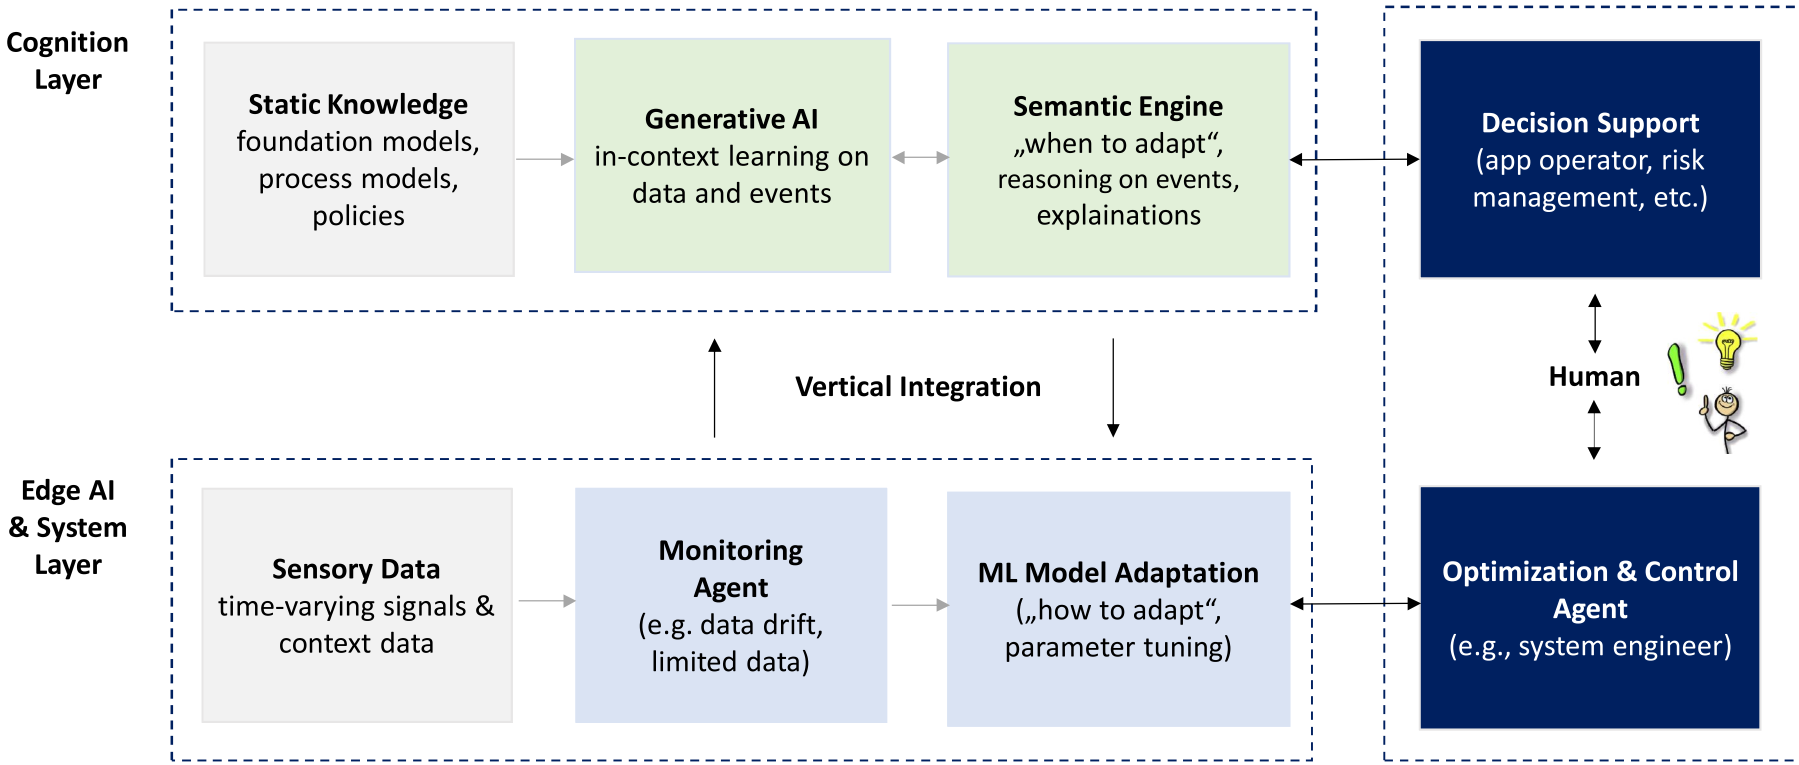

scch {
software
competence
center
hagenberg
}

Questions ...