Computer Vision and Machine Learning

IST AUSTRIA
Institute of Science and Technology

About us...

Christoph

 Asya (2012)	 Alex Z (2013)	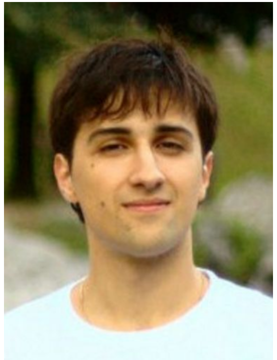 Alex K (2013)
 Amélie (2015)	 Georg (IST Fellow)	you?

About us…

central office building, 3rd floor

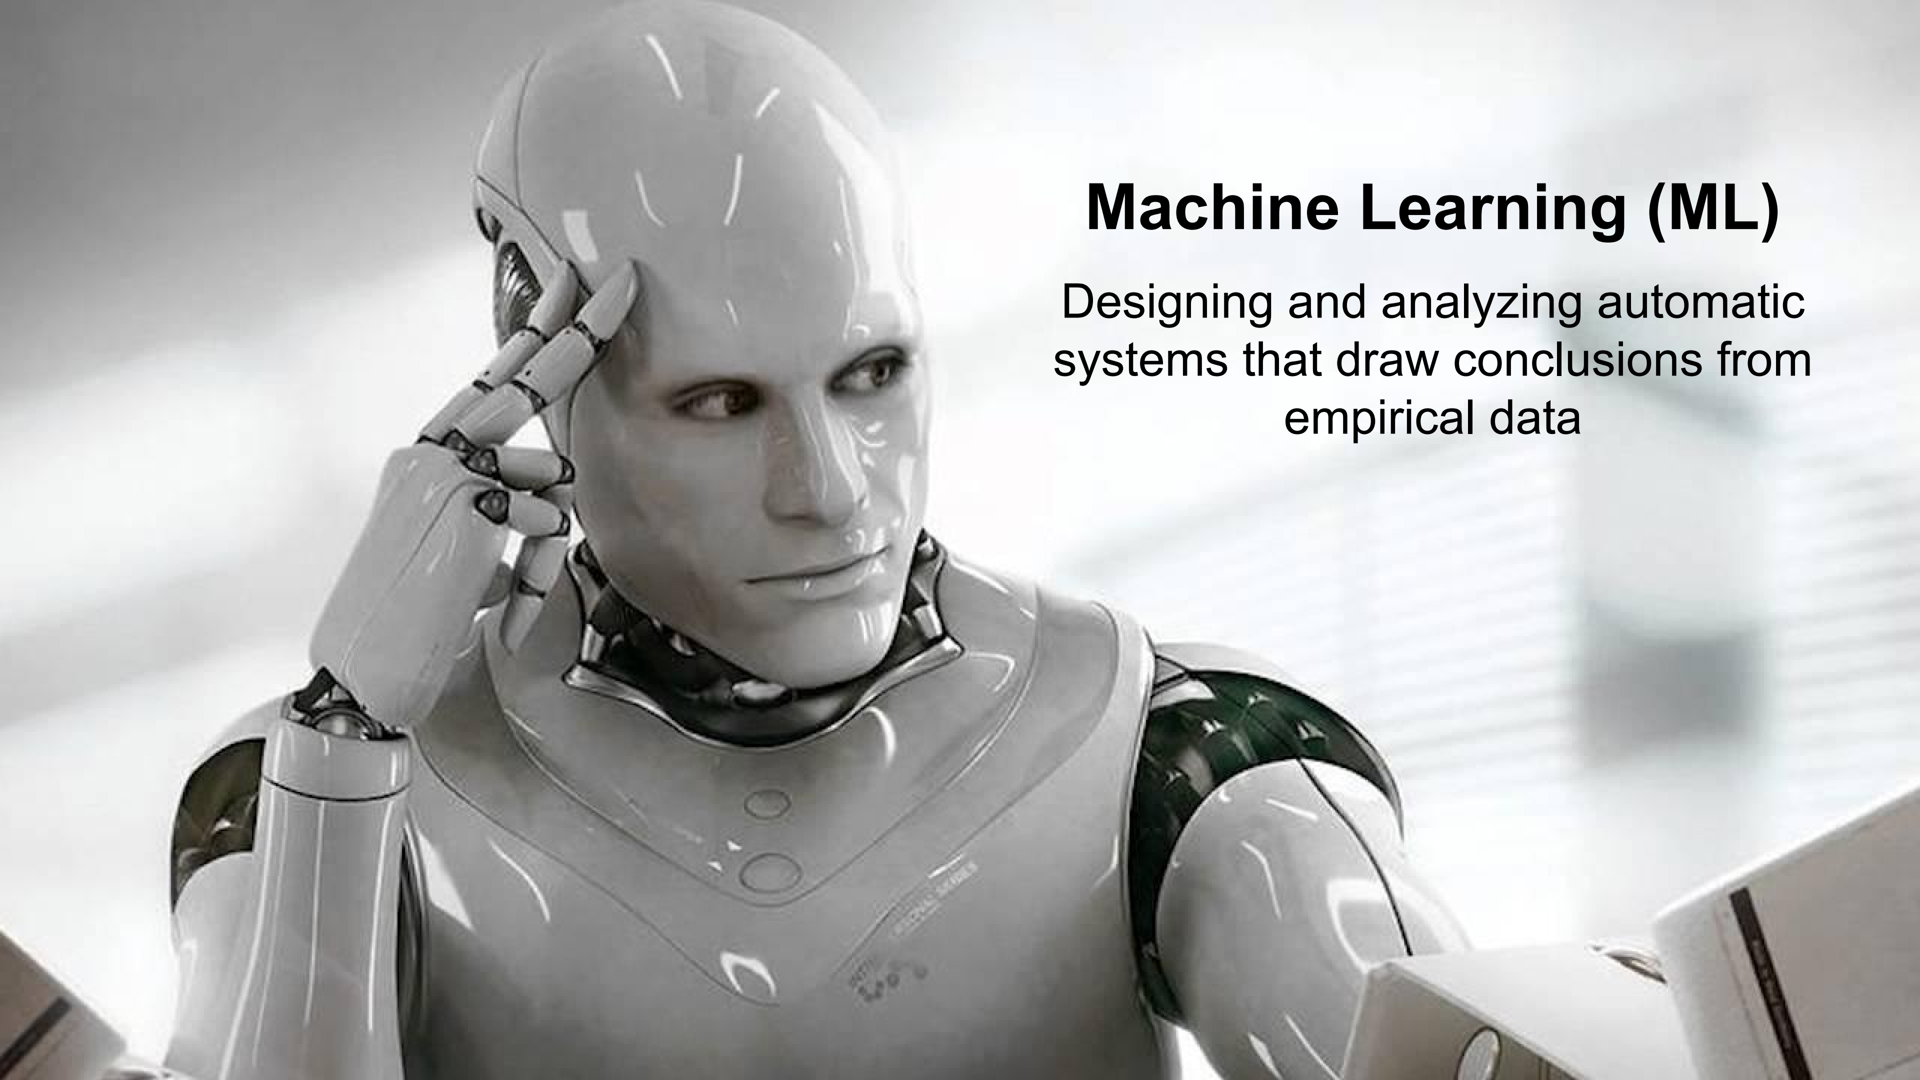

Machine Learning (ML)

Designing and analyzing automatic systems that draw conclusions from empirical data

Computer Vision (CV)

Designing and analyzing automatic systems that autonomously process visual data

"Three men sit at a table in a pub, drinking beer. One of them talks while the other two listen."

Image: British Broadcasting Corporation (BBC)

What we do

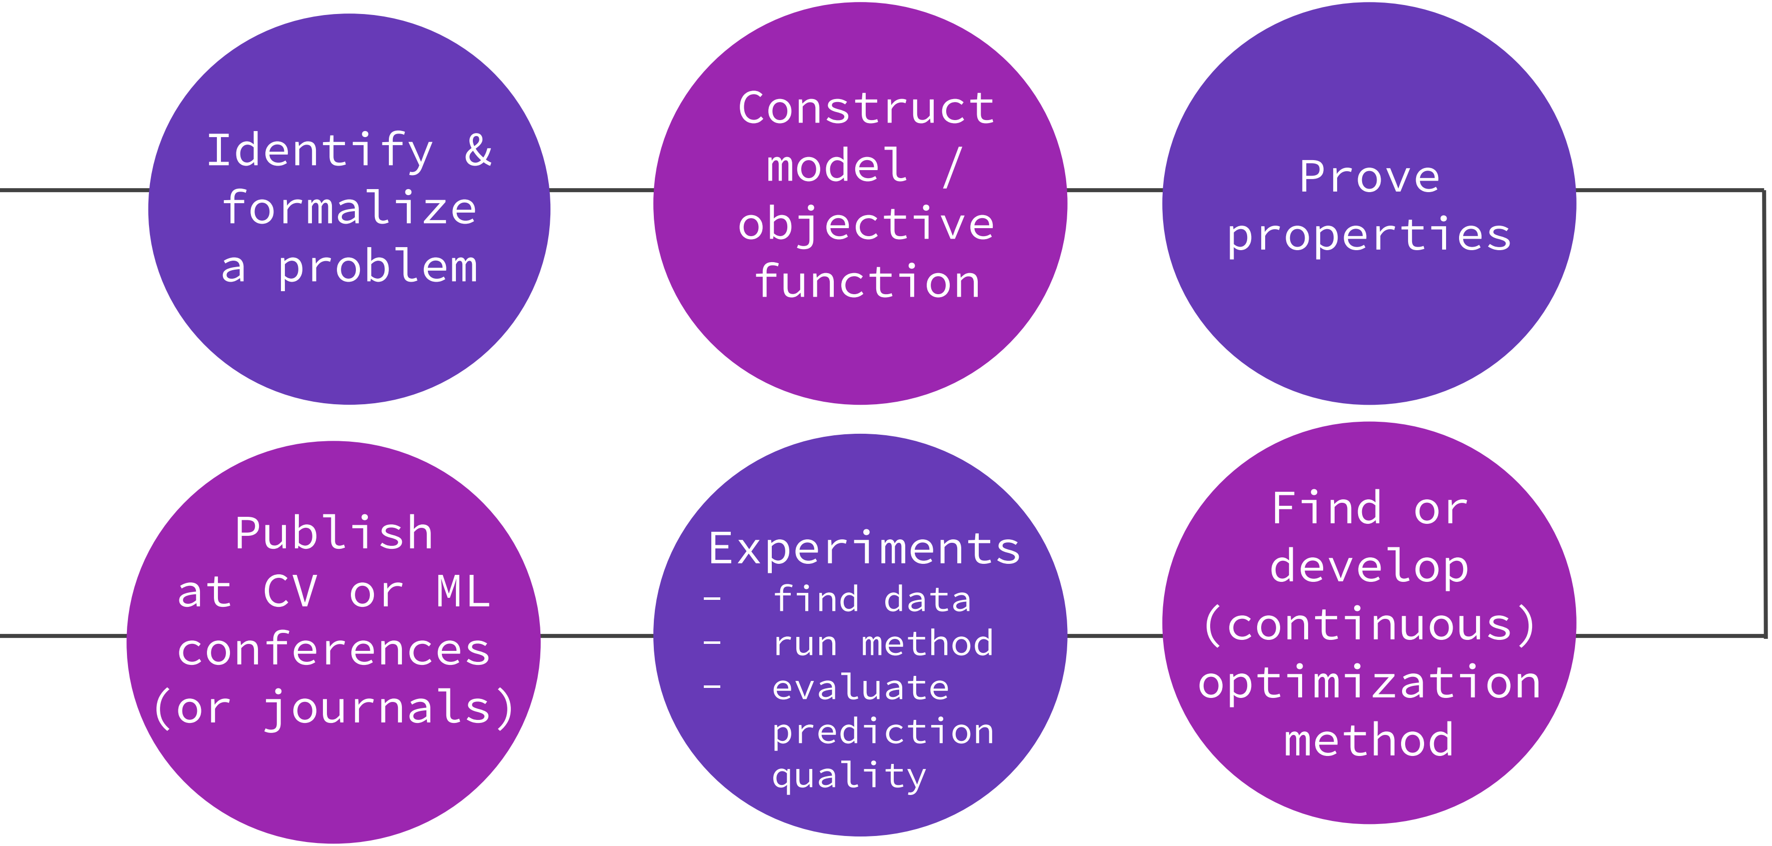
Identify & formalize a problem
Construct model / objective function
Prove properties
Find or develop (continuous) optimization method
Experiments
- find data
- run method
- evaluate prediction quality
Publish at CV or ML conferences (or journals)

Examples

identify
a problem

Extrapolation and learning equations

(Georg, CHL, in preparation for *ICLR 2016*)

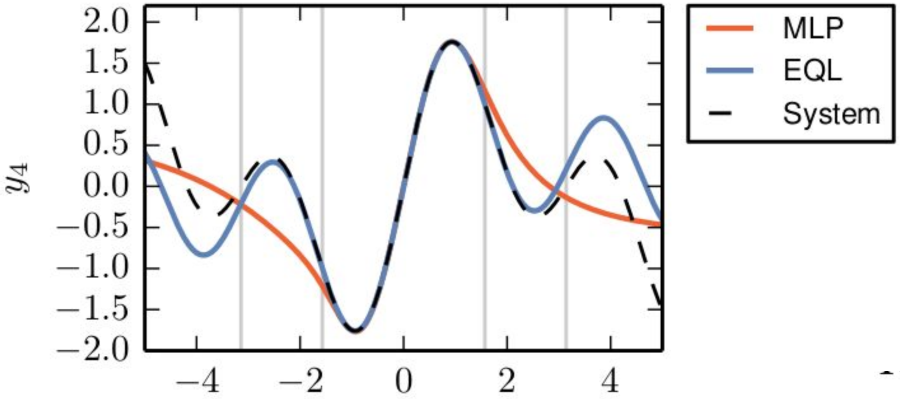

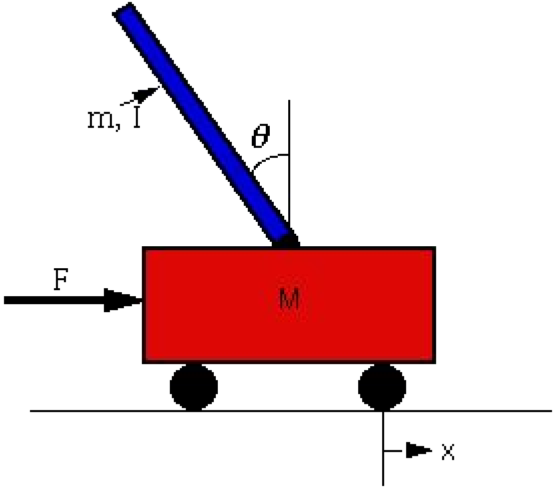

Regression methods typically find functions that **interpolate** well between observed values.
Can we learn systems that **extrapolate** well, e.g. by identifying underlying physical equations?

construct a model

Seed, Expand and Constrain: Three Principles for Weakly-Supervised Image Segmentation

(Alex K, CHL, *ECCV 2016*)

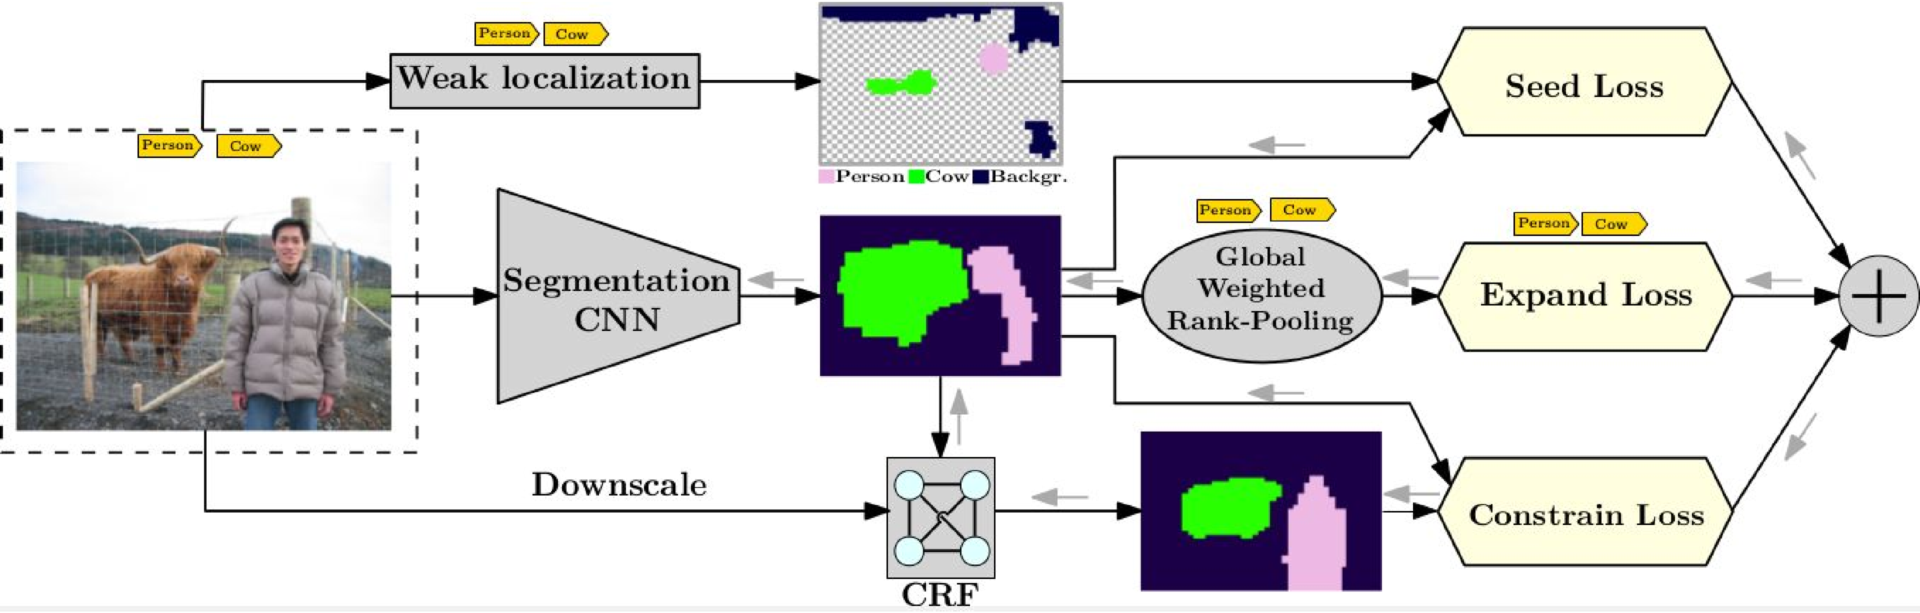

construct an objective function

Active Task Selection for Multi-Task Learning

(Asya, CHL, in preparation for *ICML 2017*)

$$\frac{1}{T}\sum_{t=1}^{T}\sum_{i\in I}\alpha_i^t \operatorname{disc}(S_t, S_i) + \frac{A}{T}\|\alpha\|_{2,1} + \frac{B}{T}\|\alpha\|_{1,2}$$

prove properties

Conditional Risk Minimization for Stochastic Processes

(Alex Z, CHL, in preparation for *AISTATS 2017*)

Theorem 2. *For any fixed n, for any $k, m \geq 1$, $\alpha \in [0,1]$ and $\beta \in [0, \alpha/4]$ the following inequality holds*

$$\mathbb{P}\left[\sup_{h\in\mathcal{H}} \left|\hat{R}_n(h) - R_n(h)\right| - \Delta_n - \Lambda_n \geq \alpha\right] \qquad (13)$$
$$\leq \frac{2k\mathcal{N}_\infty(\mathcal{L}(\mathcal{H}), \beta, n)}{(\alpha - 4\beta)^2} e^{-\frac{1}{2}m(\alpha-4\beta)^2} + \mathbb{P}\left[E^c_{k,m}\right],$$

where $\Lambda_n = \sum_{t=1}^{n} w_t(J_n) M_{t-1,J_n}$ and $\mathcal{N}_\infty(\mathcal{L}(\mathcal{H}), \beta, n)$ is a maximal β-cover of $\mathcal{L}(\mathcal{H})$ with respect to the ℓ_∞ norm (the definition is given in the appendix).

find or develop (continuous) optimization method

Multi-Plane Block-Coordinate Frank-Wolfe Algorithm for Training Structural SVMs with a Costly max-Oracle (Neel, Vladimir, CHL, *CVPR 2015*)

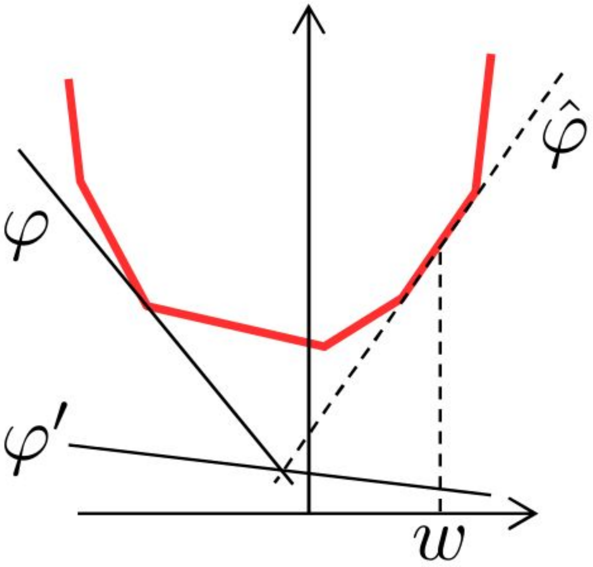

Algorithm 1 Frank-Wolfe algorithm for the dual of (4)

set $\varphi \leftarrow \varphi^{\bar{y}}$ for some $\bar{y} \in \overline{\mathcal{Y}}$
repeat
  compute $w \leftarrow \arg\min_w \frac{\lambda}{2}||w||^2 + \langle \varphi, [w\ 1] \rangle$;
  the solution is given by $w = -\frac{1}{\lambda}\varphi_\star$
  call oracle for vector w: compute $\hat{\varphi} \leftarrow \arg\max\limits_{\varphi^{\bar{y}}:\bar{y} \in \overline{\mathcal{Y}}} \langle \varphi^{\bar{y}}, [w\ 1] \rangle$
  compute $\gamma \leftarrow \arg\max_{\gamma \in [0,1]} \mathcal{F}((1-\gamma)\varphi + \gamma\hat{\varphi})$ as follows:
    set $\gamma \leftarrow \frac{\langle \varphi_\star - \hat{\varphi}_\star, \varphi_\star \rangle - \lambda(\varphi_\circ - \hat{\varphi}_\circ)}{||\varphi_\star - \hat{\varphi}_\star||^2}$ and clip γ to $[0, 1]$
  set $\varphi \leftarrow (1-\gamma)\varphi + \gamma\hat{\varphi}$
until some stopping criterion

experiments
- find data
- run method
- evaluate prediction quality

Improving Weakly-Supervised Object Localization By Micro-Annotation

(Alex K, CHL, *BMVC 2016*)

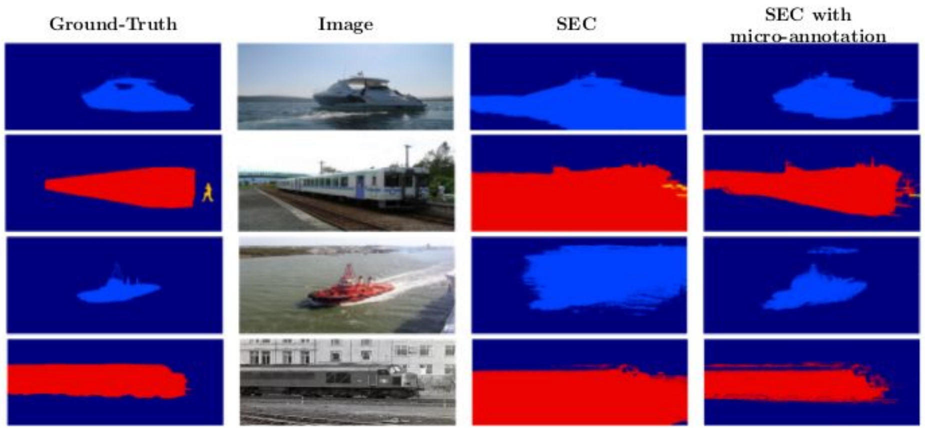

qualitative

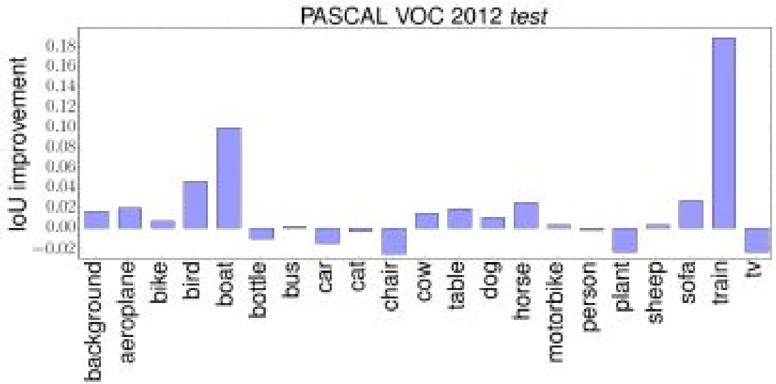

quantitative

Publish at CV or ML conferences (or journals)

Conferences (double blind, peer-reviewed, prestigious):

- Neural Information Processing Systems (NIPS), yearly
- International Conference on Machine Learning (ICML), yearly
- IEEE Computer Vision and Pattern Recognition (CVPR), yearly
- International Conference on Computer Vision (ICCV), ‘odd’ years
- European Conference on Computer Vision (ECCV), ‘even’ years

Journals:

- Journal of Machine Learning Research (JMLR)
- IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)
- International Journal of Computer Vision (IJCV)
- Machine Learning (ML)

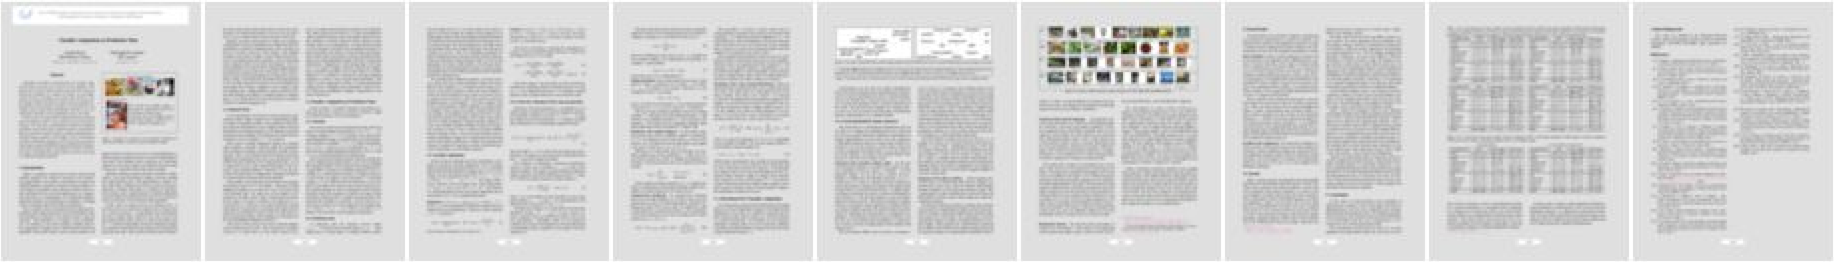

Amelie, CHL, “Classifier Adaptation at prediction time”, CVPR 2015]

Concepts we frequently use

probability
random variables, expectations, Bayes' rule, inequalities

linear algebra / calculus
function spaces, inner products, gradients, convexity

numerics/ continuous optimization
gradient-based, stochastic

public data sources
images or text, downloaded from the web

Concepts we rarely use

classical statistics
hypothesis tests, parametric data distributions

physical intuition
differential equations, dynamical system

sampling
Markov chain Monte Carlo, etc.

involved algorithms

Potential Rotation Topics

If you consider affiliating with my group	If you do not consider affiliating with my group
A topic that • what PhD research in our group is like, • builds on your prior knowledge, • ideally is useful for your actual PhD topic. Examples: • *"Metric learning for face recognition"* • *"Fisher kernels for hidden Markov models"* • *"Online guarantees for lifelong learning"*	A topic that • provides insight into CV/ML research • builds on your prior knowledge, • ideally is useful for your actual PhD topic. Examples: • Biology: *"Image processing for ant tracking"* • Cryptography: *"Learning with encrypted data"*, • Computer Graphics: *"Segmenting Meshes"*

Prerequisites

- **Mathematics**: Probability, Linear Algebra, Calculus
- **Computer Science**: Programming, preferably in Python and/or C++ (except for "theory" rotations)

Expected Outcomes

- **Presentation** in our "tea talk" series (15 minutes)
- **Written report** (5 to 10 pages)

Recommended Courses

- **Core course:** guest lecture *"Clustering"*
- **Track core courses:** *"Data Science and Scientific Computing"* or *"Computer Science"*
- **Autumn 1:** *"Methods of Data Analysis" (G. Tkacik)*
- **Autumn 2:** *"Probabilistic graphical model"* (CHL)
- **Spring 1:** *"Numerical Algorithms" (C. Wojtan)*
- **Spring 2:** *"Applications of Stochastic Processes" (N.Barton)*

Group Events

- Tuesdays 10:45 **"Tea talks"** (15 min. talk series)
- Tuesdays 11:00 **CVML Reading group**

Office Hours

- **open door**
- or send me email: `chl@ist.ac.at`