AAAI-17 San Francisco

AAAI-17 Review

Lantao Yu

Mar. 8 2017

AAAI Attendance Trend

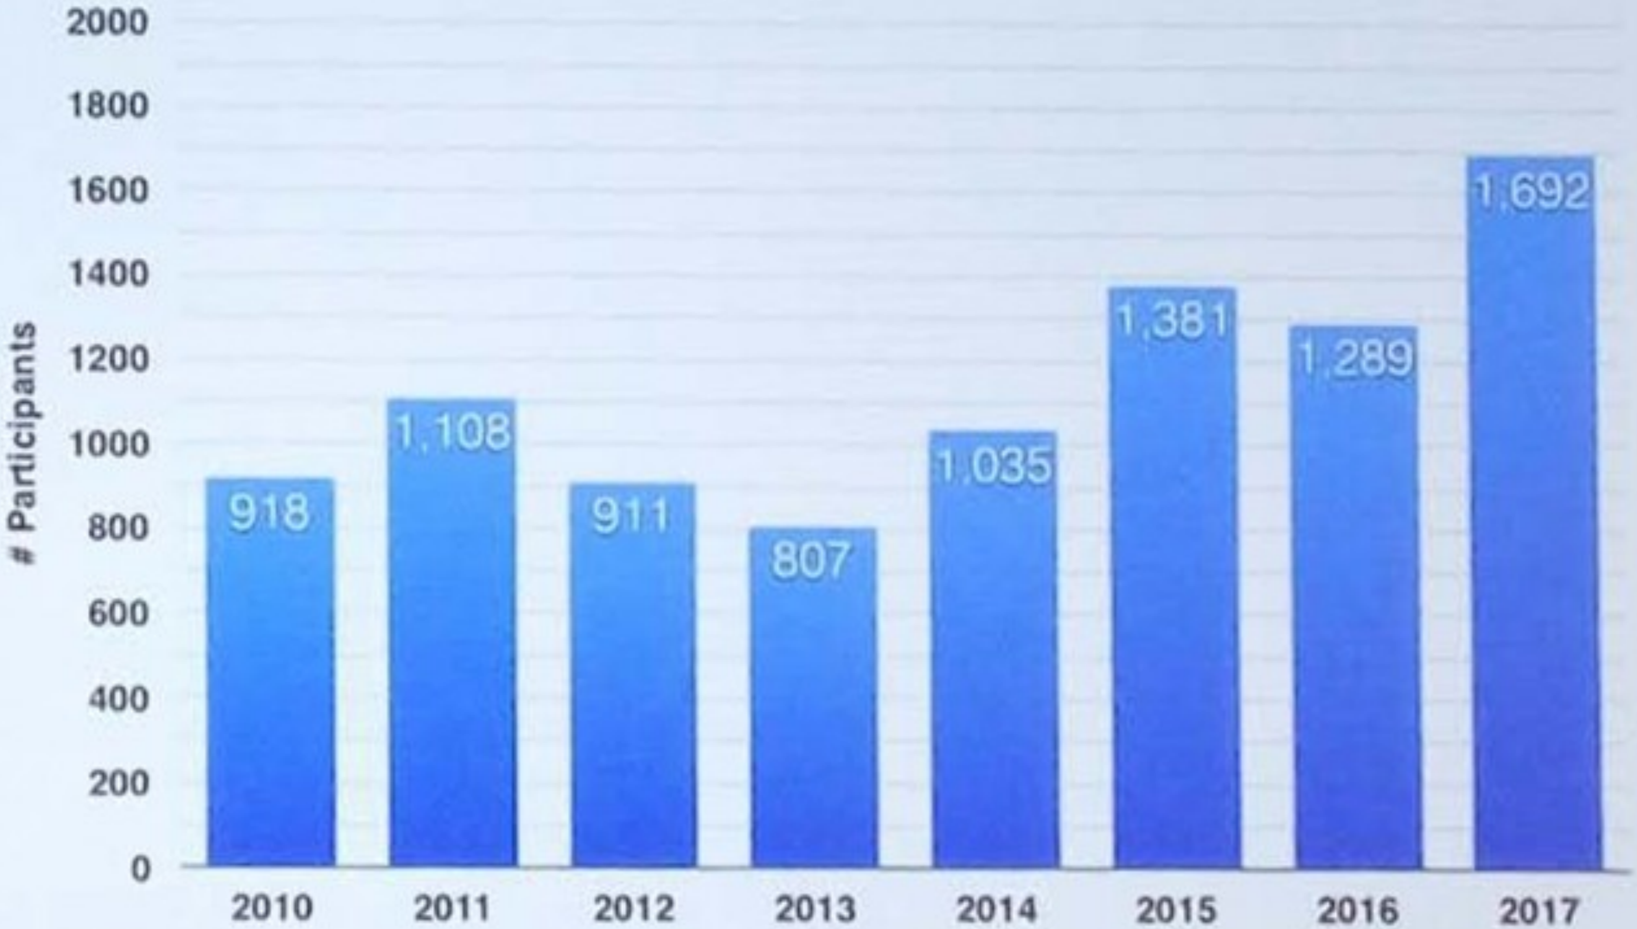

AAAI Submission Trend

A paper with 3 authors from country c1 and 1 from country c2 is considered 0.75 c1 and 0.25 c2

★★★

Number of Submissions per Country

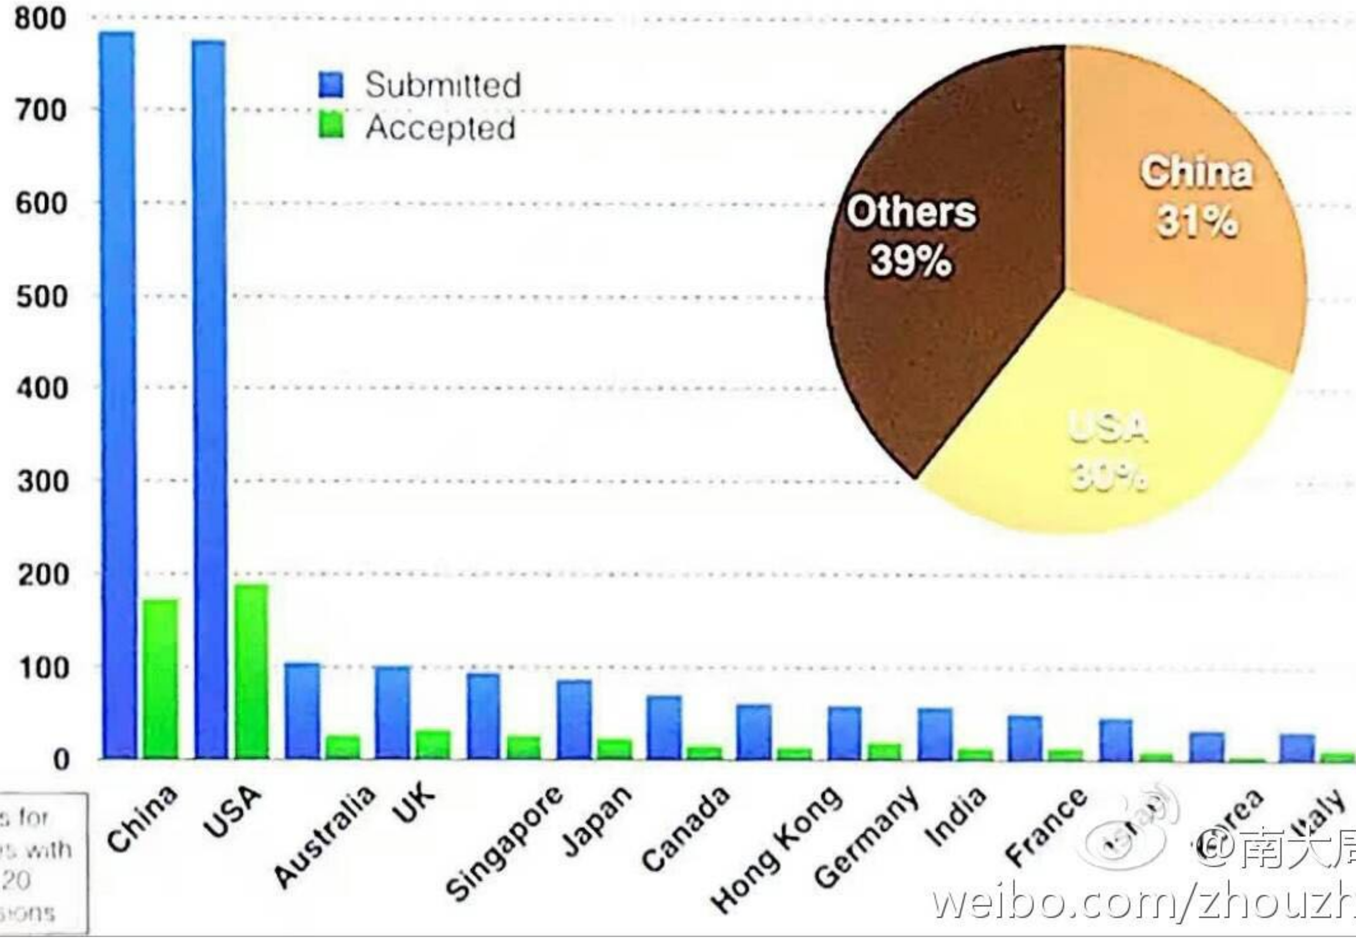

statistics for countries with at least 20 submissions

	Submitted		Accepted	
Machine Learning	910	0.35	215	0.24
NLP	373	0.15	77	0.21
Applications	268	0.10	61	0.23
Search/planning	194	0.08	59	0.31
Vision	169	0.07	45	0.27
Knowledge Rep	131	0.05	33	0.25
Game Theory	113	0.04	43	0.38
MAS	89	0.03	21	0.24
Uncertainty	75	0.03	20	0.28
Human	66	0.03	14	0.22
Cognitive Systems Track	58	0.02	16	0.28
Comp Sustainability Track	54	0.02	14	0.26
Intelligent Systems Track	27	0.01	8	0.30
Robotics	26	0.01	6	0.23

***

	Submitted	Accepted	Rate	PCs	SPCs
Main Technical Track	2432	601	0.247	1349	168
Computational Sustainability	54	14	0.259	32	
Cognitive Systems	58	16	0.276	52	
Integrated System	27	8	0.296	35	
Total of 4 technical tracks	2571	639	0.249	1468	168

	Submitted	Accepted	Rate	PCs
Demo Program	35	13	0.371	8
Senior Member	22	12	0.545	30
Doctoral Consortium	31	16	0.516	31
What's Hot Track		10		
Student Abstracts	115	67	0.583	44

10 most frequent words in paper abstracts

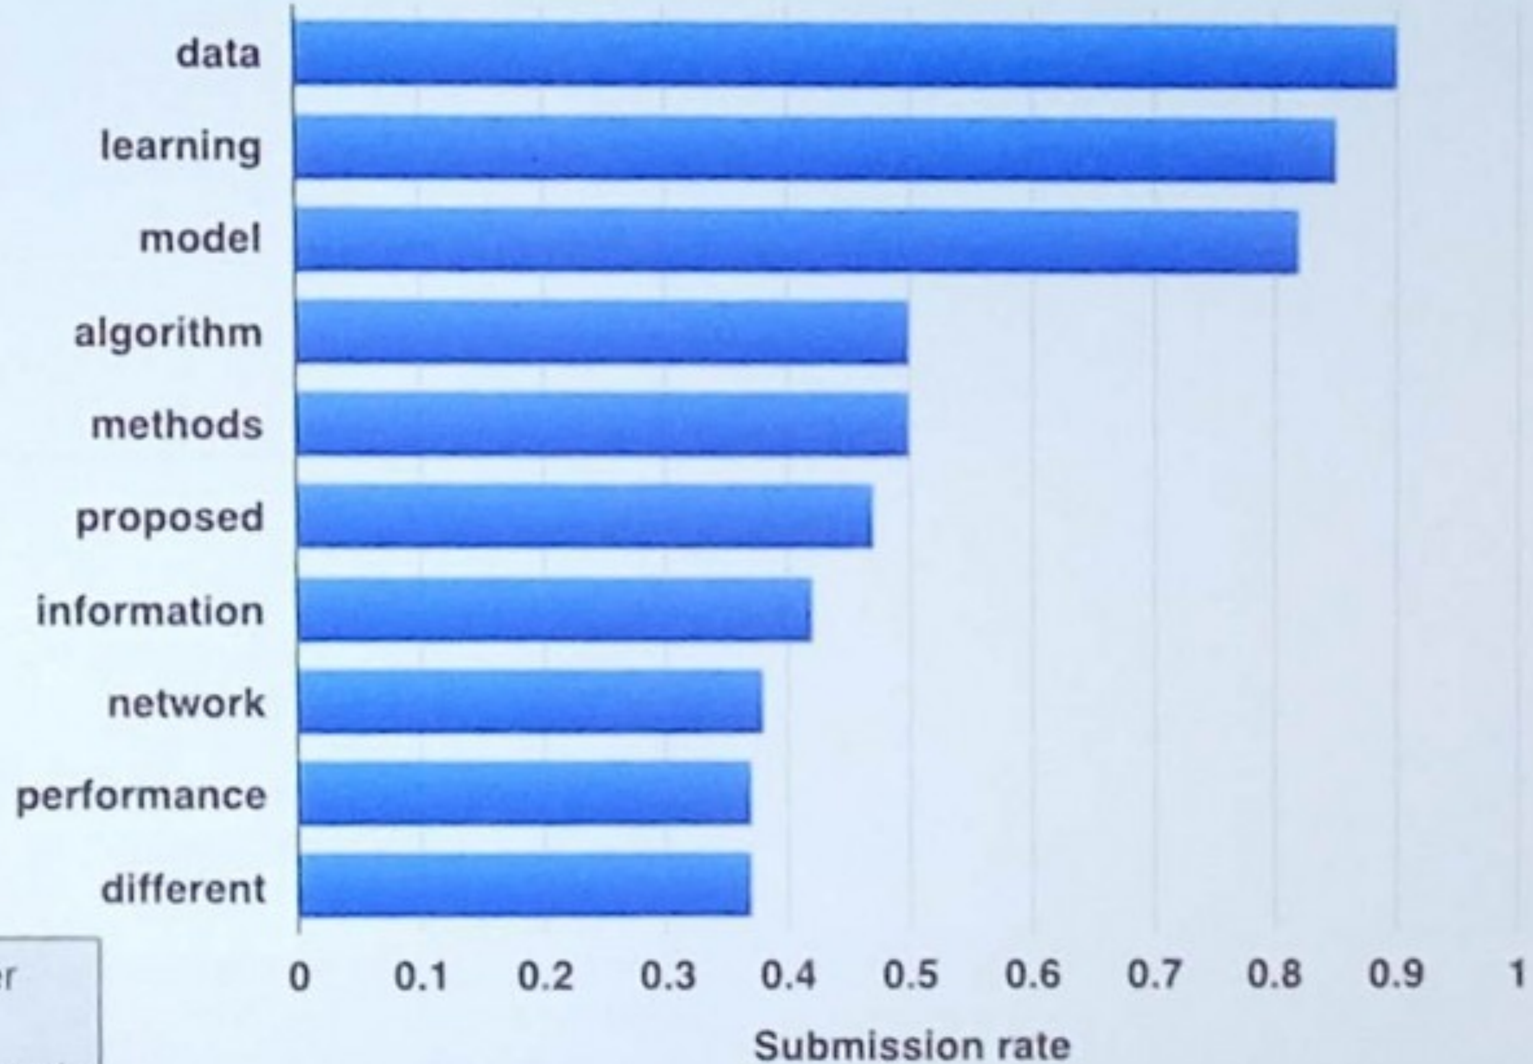

statistics over words that appear at least in 20 papers

Acceptance Rate by Area

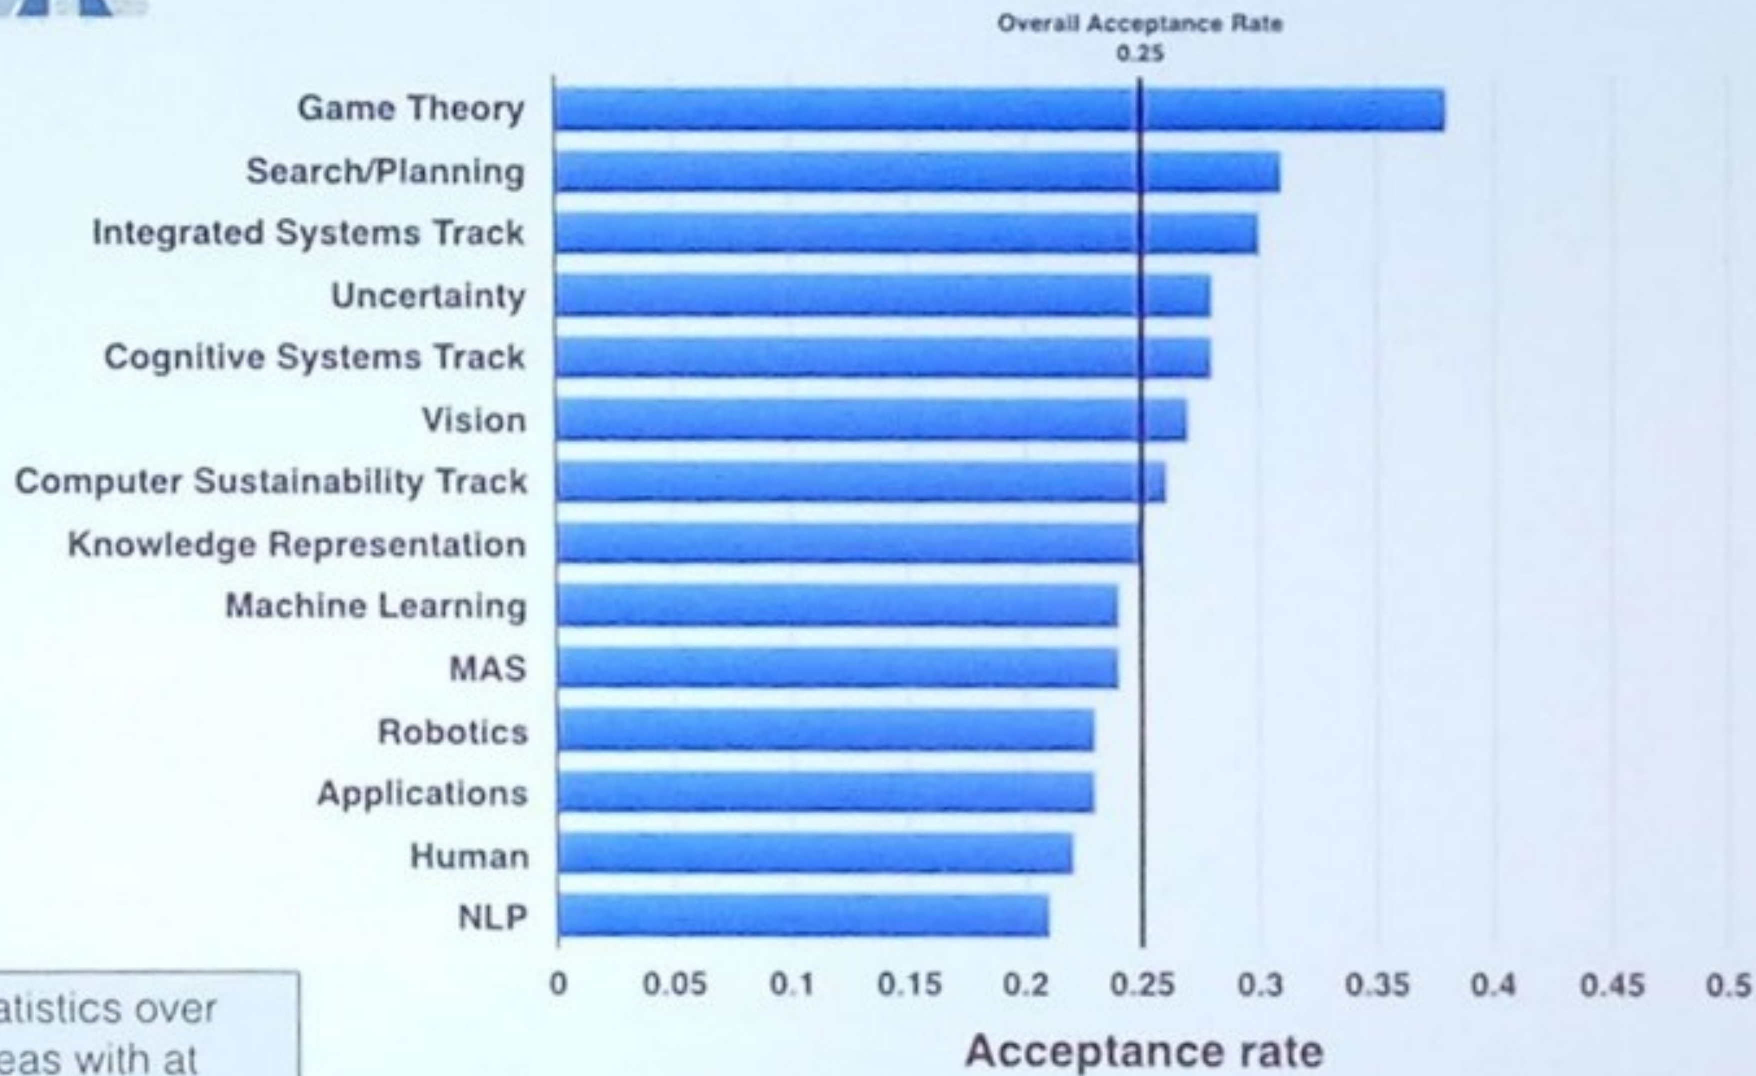

statistics over areas with at least in 20 papers

AAAI-17 Outstanding Paper Award

Label-Free Supervision of Neural Networks with Physics and Domain Knowledge

Russell Stewart , Stefano Ermon
Department of Computer Science, Stanford University

Traditional Supervised Learning Setting

$$D = \{(x_1, y_1), \ldots, (x_n, y_n)\}$$

The goal is to learn a function $f : X \rightarrow Y$ mapping inputs to outputs.

The goal of our method is to train a network, f, mapping from inputs to outputs that we care about, without needing direct examples of those outputs.

The goal of our method is to train a network, f, mapping from inputs to outputs that we care about, without needing direct examples of those outputs.

$$f^* = \arg\min_{f \in \mathcal{F}} \sum_{i=1}^{n} \ell(f(x_i), y_i)$$

$$\hat{f}^* = \arg\min_{f \in \mathcal{F}} \sum_{i=1}^{n} g(x_i, f(x_i)) + R(f)$$

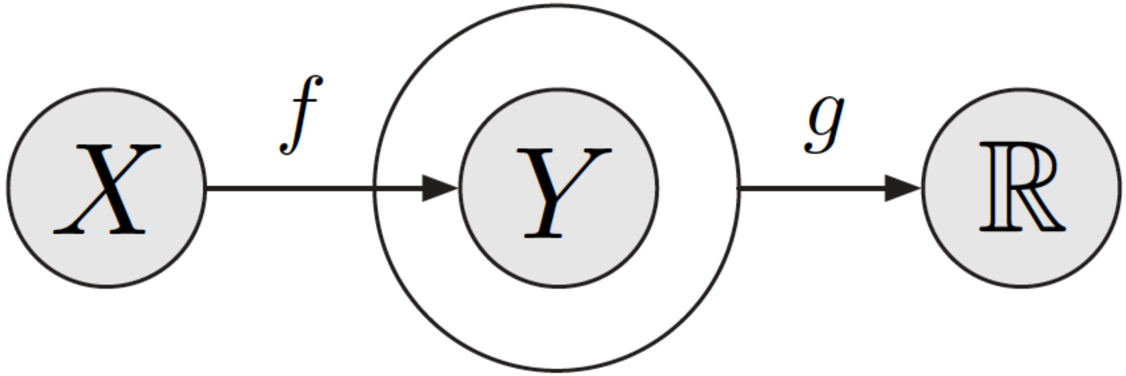

Figure 1: Constraint learning aims to recover the transformation f without providing labels y. Instead, we look for a mapping f that captures the structure required by g.

$$\hat{f}^* = \arg\min_{f \in \mathcal{F}} \sum_{i=1}^{n} g(x_i, f(x_i)) + R(f)$$

Tracking an object in free fall

- In the first experiment, we record videos of an object being thrown across the field of view and aim to learn the object's height in each frame.
- Mapping: $\left(\mathbb{R}^{\text{height}\times\text{width}\times 3}\right)^N \rightarrow \mathbb{R}^N$
- the plot of the object's height over time will form a parabola:

$$\mathbf{y}_i = y_0 + v_0(i\Delta t) + a(i\Delta t)^2$$

Tracking an object in free fall

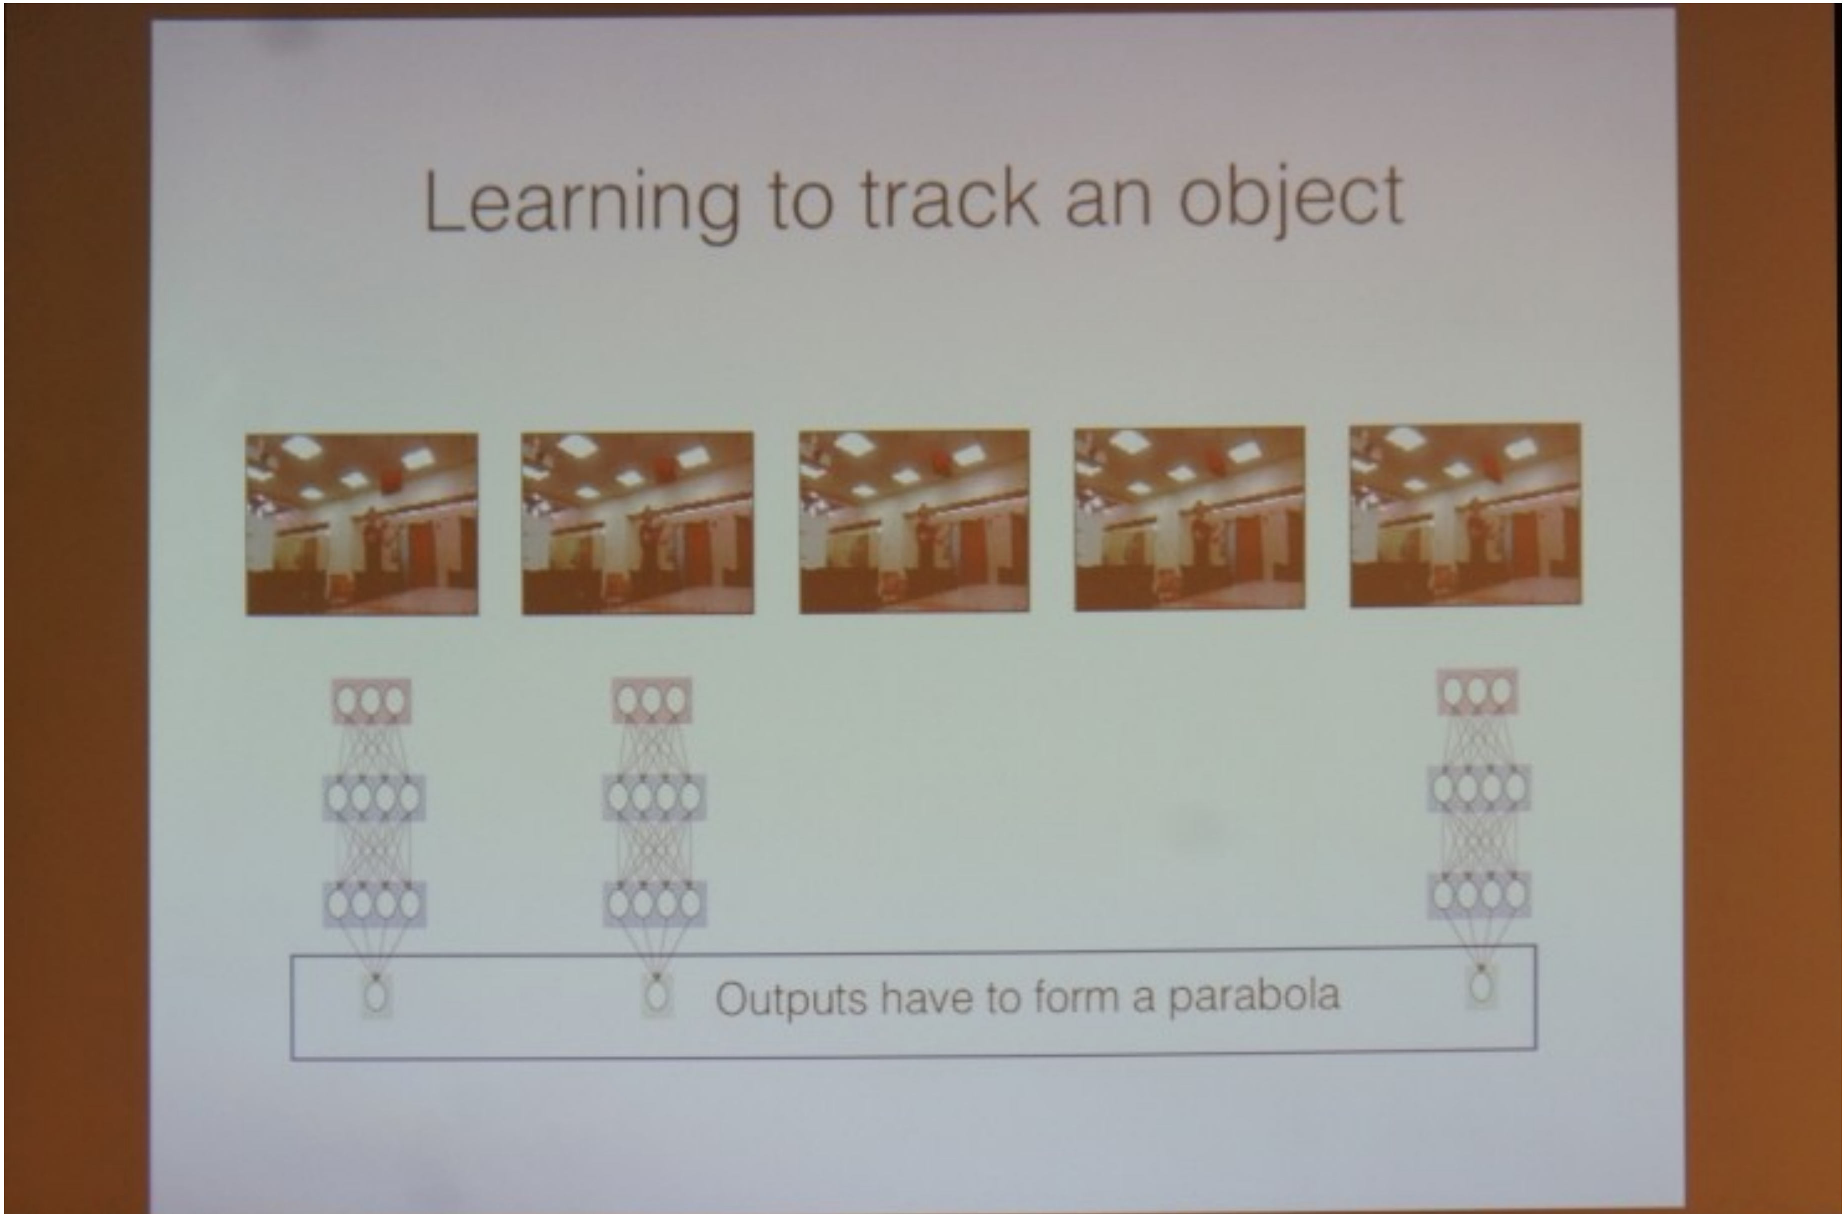

Given any trajectory of N height predictions, f(x) , we fit a parabola with fixed curvature to those predictions, and minimize the resulting residual.

$$\boldsymbol{a} = [a\Delta t^2, a(2\Delta t)^2, ..., a(N\Delta t)^2]$$

$$\hat{\mathbf{y}} = \mathbf{a} + \mathbf{A}(\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T(f(\mathbf{x}) - \mathbf{a}) \tag{3}$$

where

$$\mathbf{A} = \begin{bmatrix} \Delta t & 1 \\ 2\Delta t & 1 \\ 3\Delta t & 1 \\ \vdots & \vdots \\ N\Delta t & 1 \end{bmatrix}$$

$$g(\mathbf{x}, f(\mathbf{x})) = g(f(\mathbf{x})) = \sum_{i=1}^{N} |\hat{\mathbf{y}}_i - f(\mathbf{x})_i|$$

Tracking an object in free fall

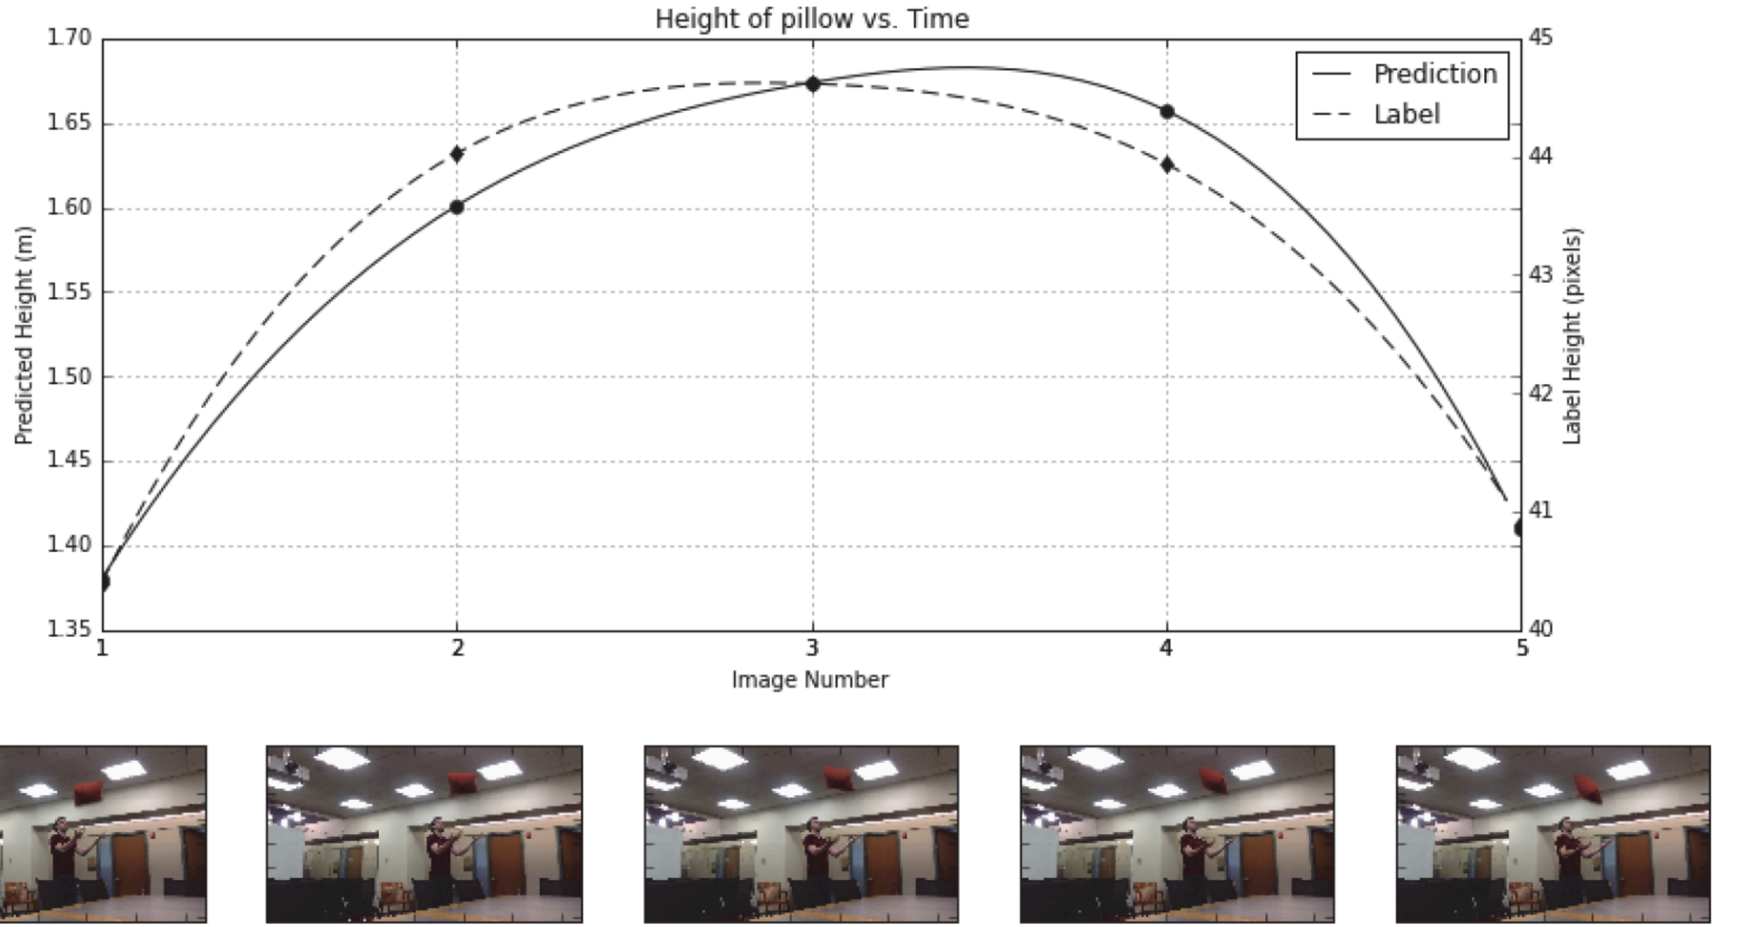

Label-free 90.1% vs 94.5 Supervised Learning

Tracking the position of a walking man

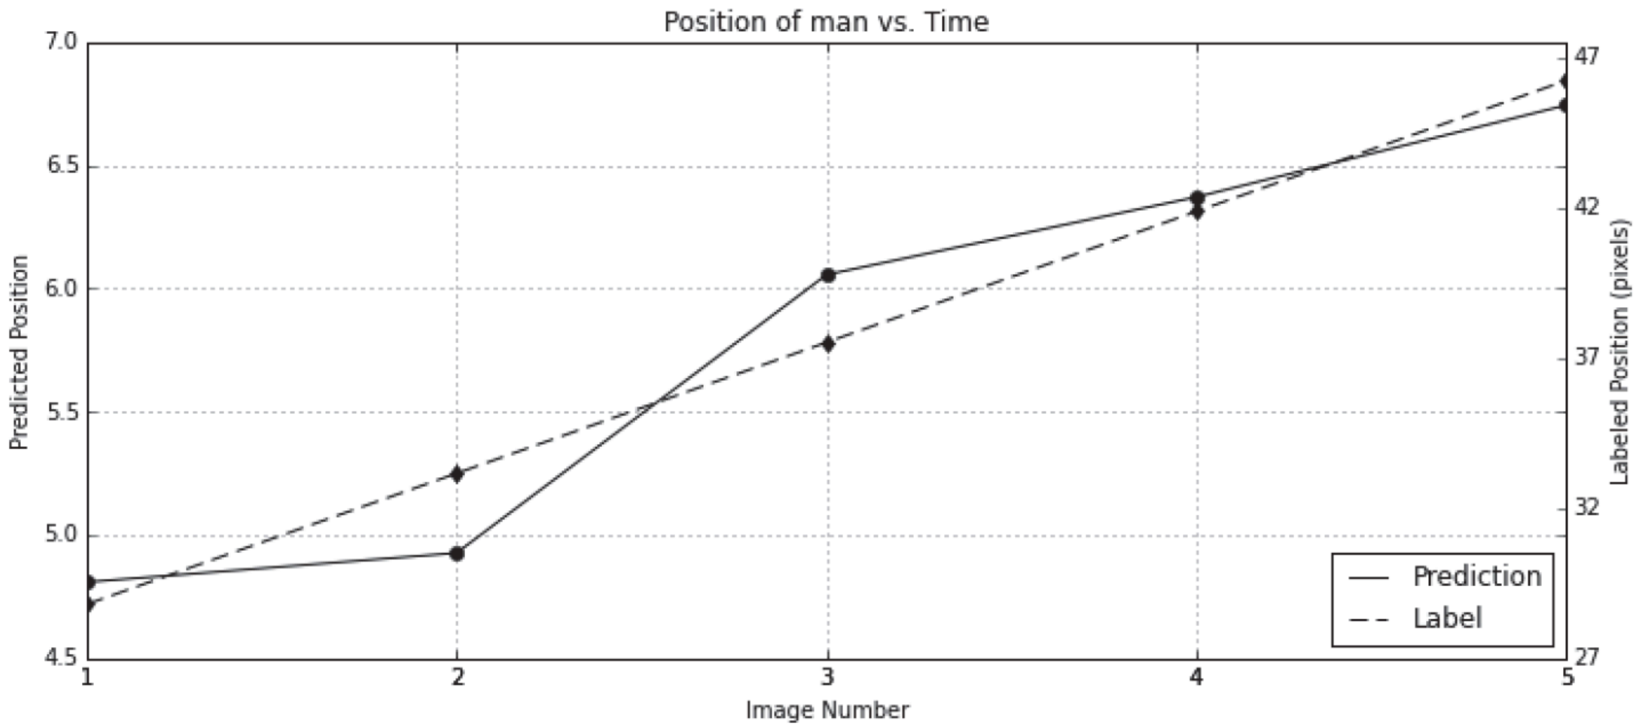

Detecting objects with causal relationships

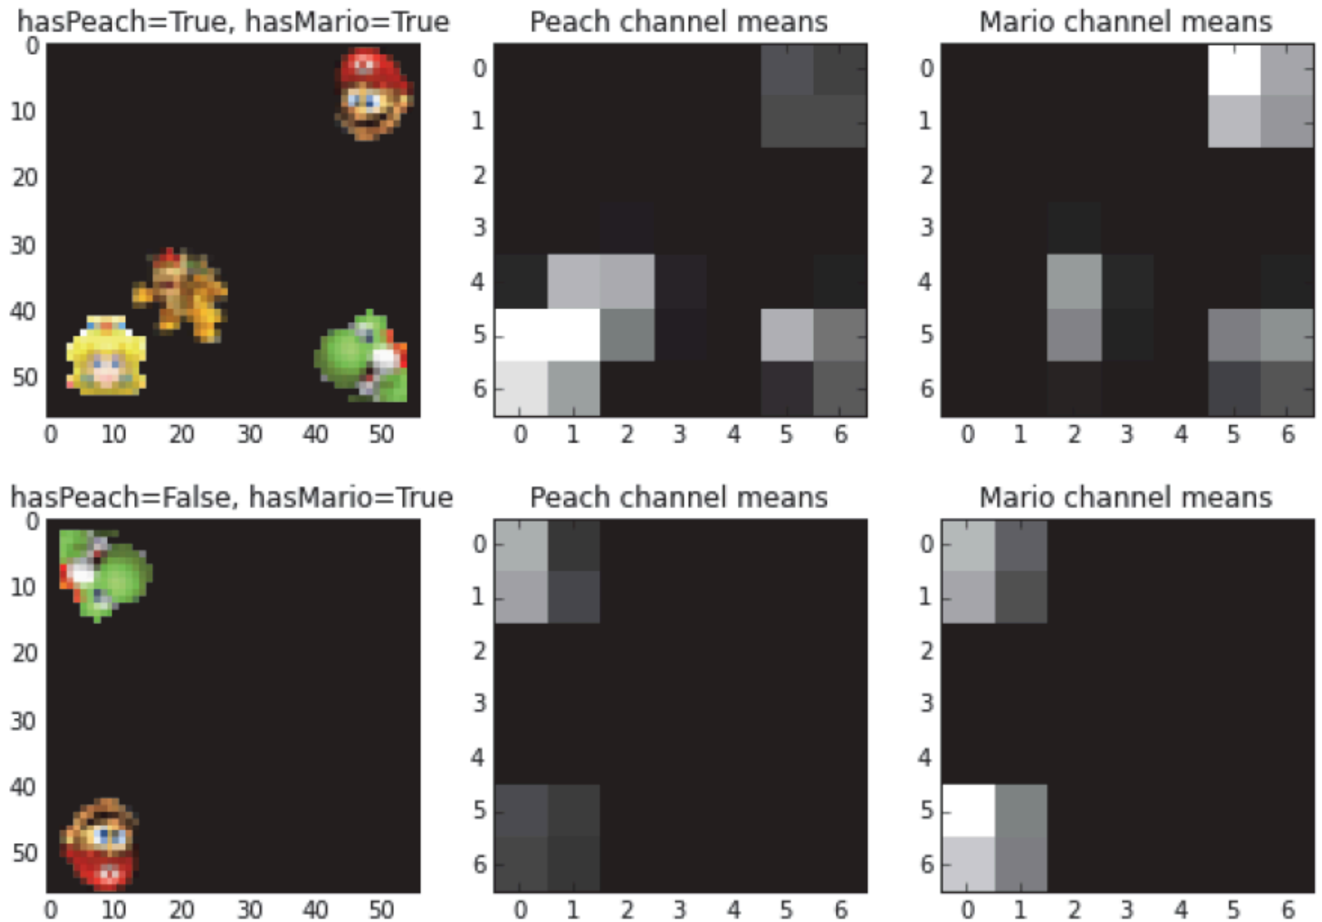

Task: predict whether each image contains Mario (red) and Peach (yellow), knowing only that Peach => Mario

Detecting objects with causal relationships

- Explore the possibilities of learning from logical constraints imposed on single images.
- Our aim is to create a pair of neural networks f = (f1 , f2) for identifying Peach and Mario, respectively.
- Rather than supervising with direct labels, we train the networks by constraining their outputs to have the logical relationship y1 $\Rightarrow$ y2
- Need three complicated loss function:
    - h1 forces rotational independence of the output by applying a random horizontal and vertical reflection ρ, to images.
    - h2 and h3 allows us to avoid trivial solutions by encouraging high standard deviation and high entropy outputs, respectively.

$$h_1(\mathbf{x}, k) = \frac{1}{M}\sum_{i}^{M} |Pr[f_k(\mathbf{x}) = 1] - Pr[f_k(\rho(\mathbf{x})) = 1]|$$

$$h_2(\mathbf{x}, k) = - \operatorname*{std}_{i \in [1 \ldots M]} (Pr[f_k(\mathbf{x}_i) = 1])$$

$$h_3(\mathbf{x}, v) = \frac{1}{M}\sum_{i}^{M} (Pr[f(\mathbf{x}_i) = v] - \frac{1}{3} + (\frac{1}{3} - \mu_v))^2$$

$$\mu_v = \frac{1}{M}\sum_{i}^{M} \mathbb{1}\{v = \operatorname*{arg\,max}_{v' \in \{0,1\}^2} Pr[f(\mathbf{x}) = v']\}$$

Distant Domain Transfer Learning

Ben Tan,[*] Yu Zhang,[*] Sinno Jialin Pan,[**] Qiang Yang[*]
[*]Hong Kong University of Science and Technology, Hong Kong
[**]Nanyang Technological University, Singapore
{btan,zhangyu,qyang}@cse.ust.hk, sinnopan@ntu.edu.sg

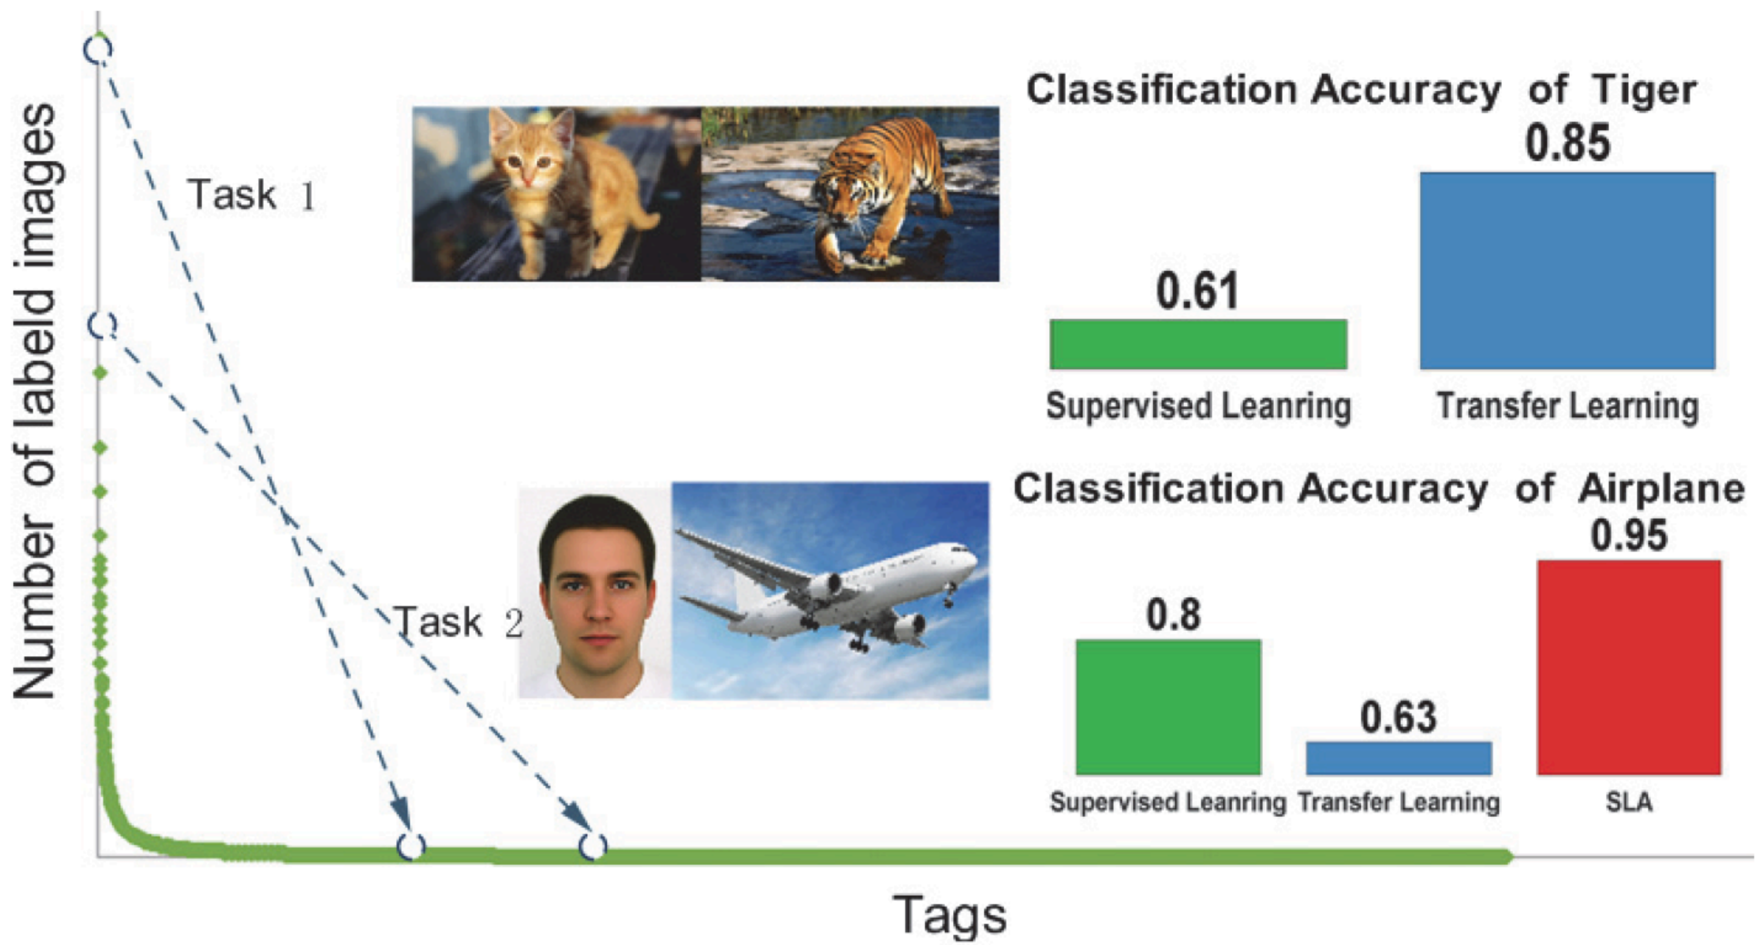
Number of labeld images
Task 1
Task 2
Classification Accuracy of Tiger
0.61
0.85
Supervised Leanring
Transfer Learning
Classification Accuracy of Airplane
0.8
0.63
0.95
Supervised Leanring
Transfer Learning
SLA
Tags

- Inspired by human’s ‘transitivity’ learning and inference ability
- From one intermediate domain => multiple intermediate domain
- Automatically select the useful subset of data from each intermediate domain
- Almost all existing transfer learning methods (instance weighting, feature mapping, model adaptation) assume the source and target domain are closely related.
- The goal of DDTL is to exploit the unlabeled data in the intermediate domains to build a bridge between the source and target domains

$$\mathcal{J}_1(f_e, f_d, \boldsymbol{v}_S, \boldsymbol{v}_T) = \frac{1}{n_S}\sum_{i=1}^{n_S} v_S^i \|\hat{\boldsymbol{x}}_S^i - \boldsymbol{x}_S^i\|_2^2 + \frac{1}{n_I}\sum_{i=1}^{n_I} v_I^i \|\hat{\boldsymbol{x}}_I^i - \boldsymbol{x}_I^i\|_2^2$$

$$+ \frac{1}{n_T}\sum_{i=1}^{n_T} \|\hat{\boldsymbol{x}}_T^i - \boldsymbol{x}_T^i\|_2^2 + R(\boldsymbol{v}_S, \boldsymbol{v}_T), \quad (1)$$

$$R(\boldsymbol{v}_S, \boldsymbol{v}_T) = -\frac{\lambda_S}{n_S}\sum_{i=1}^{n_S} v_S^i - \frac{\lambda_I}{n_I}\sum_{i=1}^{n_I} v_I^i.$$

Incorporating side information (source and target domains have label information)

$$\mathcal{J}_2(f_c, f_e, f_d) = \frac{1}{n_S}\sum_{i=1}^{n_S} v_S^i \ell(y_S^i, f_c(\boldsymbol{h}_S^i)) + \frac{1}{n_T}\sum_{i=1}^{n_T} \ell(y_T^i, f_c(\boldsymbol{h}_T^i))$$

$$+ \frac{1}{n_I}\sum_{i=1}^{n_I} v_I^i g(f_c(\boldsymbol{h}_I^i)), \quad (2)$$

Table 2: Accuracies (%) of selected tasks on Catech-256 and AwA with SIFT features.

	SVM	DTL	GFK	LAN	ASVM	TTL	STL	SLA
‘horse-to-face’	84 ± 2	88 ± 2	77 ± 3	79 ± 2	76 ± 4	78 ± 2	86 ± 3	$\mathbf{92} \pm 2$
‘airplane-to-gorilla’	75 ± 1	62 ± 3	67 ± 5	66 ± 4	51 ± 2	65 ± 2	76 ± 3	$\mathbf{84} \pm 2$
‘face-to-watch’	75 ± 7	68 ± 3	61 ± 4	63 ± 4	60 ± 5	67 ± 4	75 ± 5	$\mathbf{88} \pm 4$
‘zebra-to-collie’	71 ± 3	69 ± 2	56 ± 2	57 ± 3	59 ± 2	70 ± 3	72 ± 3	$\mathbf{76} \pm 2$

Thank you!