Learning with Less Labels (LwLL)

Wade Shen

Goal

1. Create AI/ML algorithms that can learn with 10^6+ less labels and adapt to new environments with $<10^2$ labels
2. Find the theoretical limits on learning and adaptation given tasks and data

The labeling problem

1. Need lots of labels

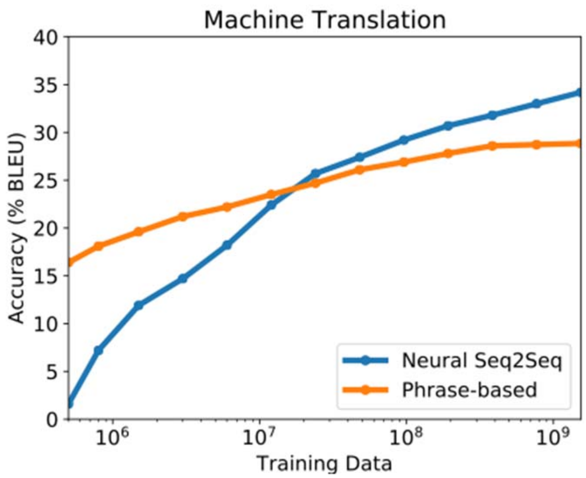

Task	Data
Object recognition	100+m objs
Speech recognition	35k+ hrs (10^{10} frames)
Translation	10^{10} words

2. Domain change = catastrophic failure

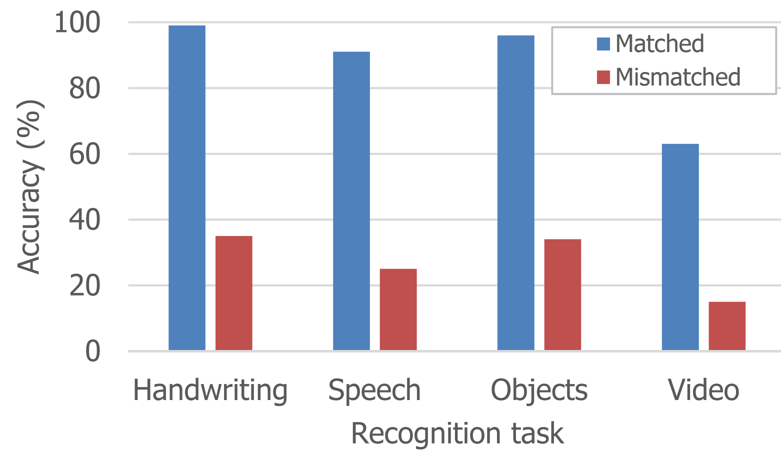

3. Labeling USG data is expensive

Crowd Sourcing

	Google	USG
Available labor	10^5-10^6	100-1000
Cost of labor	$1.25/hr	$34.00/hr
Cost to translation model	$41m	$~1bn
Time to translation model	0.5 yrs	17 yrs

4. You need them again!

Speaker ID

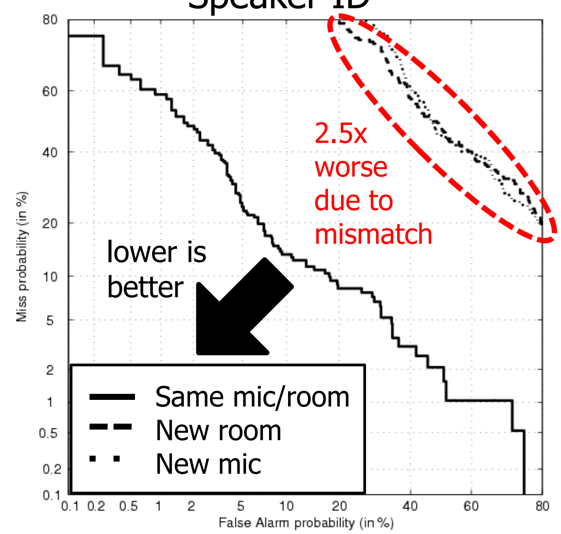

Learning with Less Labels (LwLL)

Today: Costly training and retraining

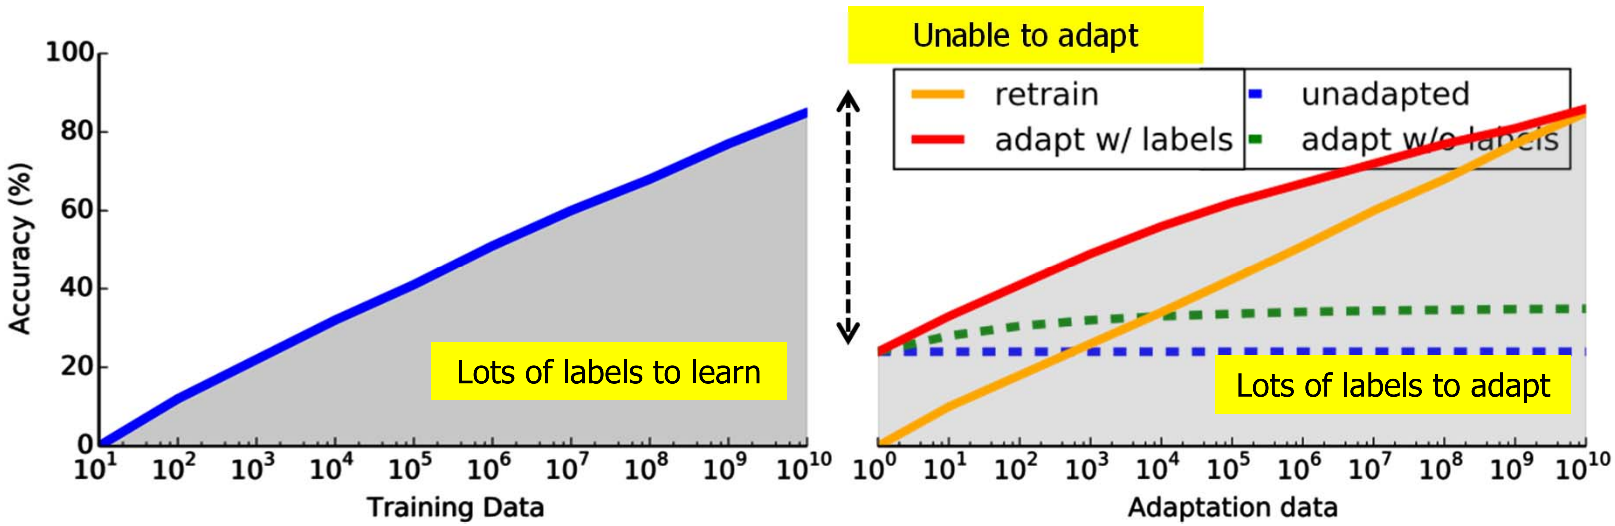

Tomorrow: Learning with less labels

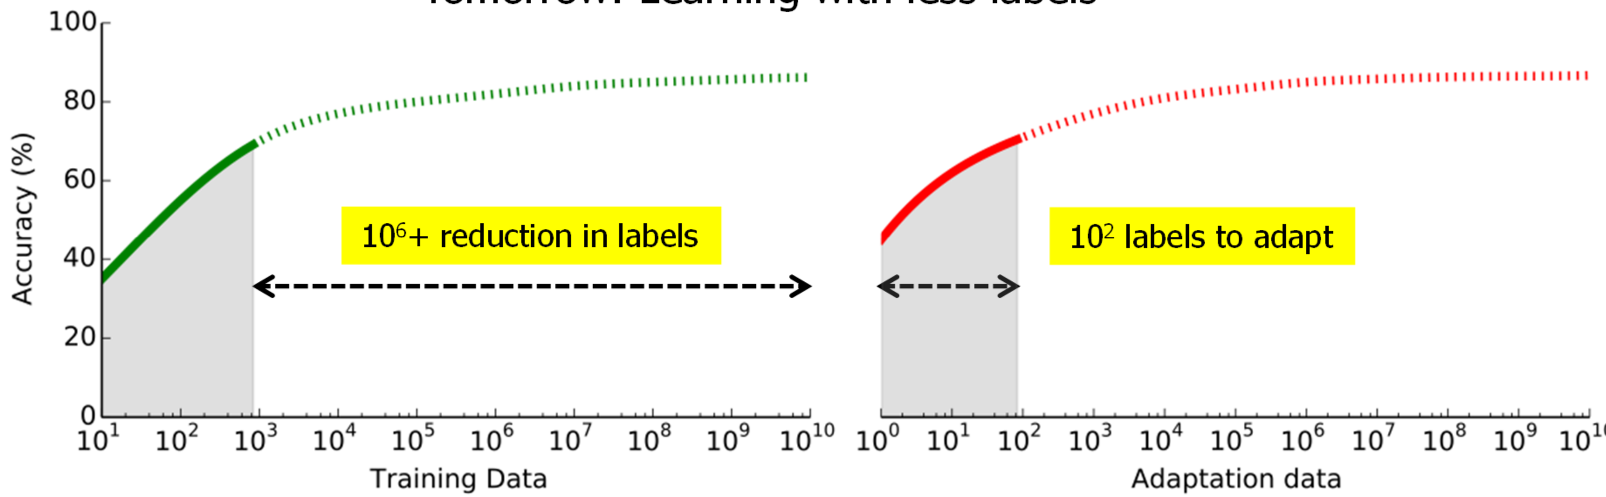

LwLL: Supervision via generalized objectives

- Learning boils down to optimization

The objective function

The model

$$\arg\min_{w} E(\mathbf{w}) = \frac{1}{N}\sum_{i=1}^{N}(\mathbf{w}\cdot\mathbf{x_i} - \mathbf{y_i})^2$$

i^{th} label

The i^{th} data observation

- Problems and opportunities:

1. Observations aren't ideal

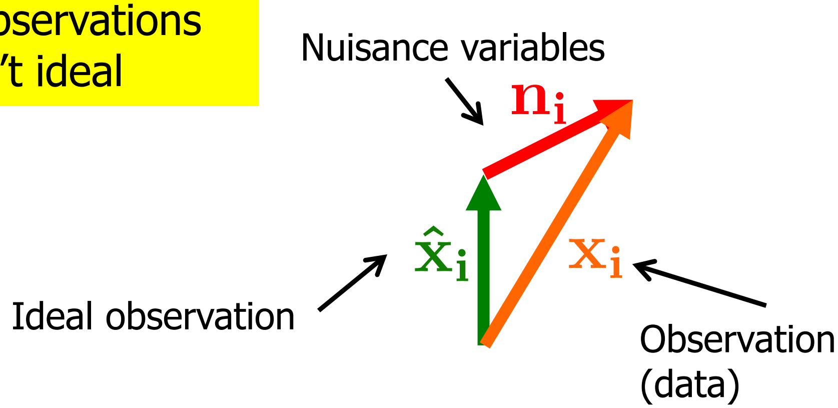

2. Learning requires data paired with labels

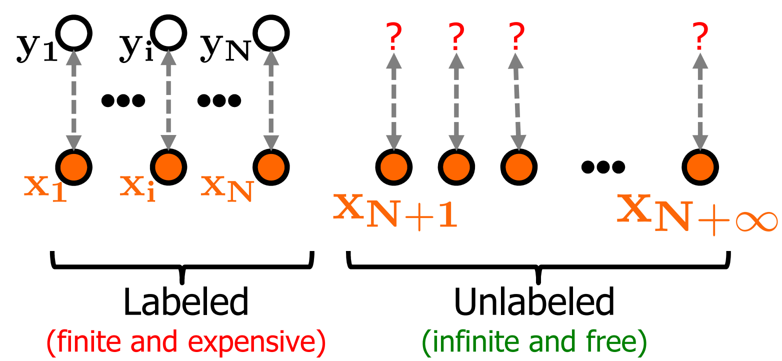

LwLL will:

1) Automatically find a generalized objective *J(**w**)* that eliminates nuisance and efficiently utilizes unlabeled data
2) Formally define the limits on *J(**w**)* for adaption and training given data and labels

TA1: Autonomously learn what matters

- **Goal:** discover the data that matters

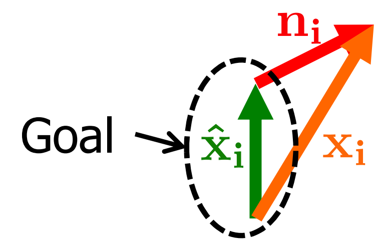

- **Approaches:**

1. **Automated** transfer learning that discovers similar problems and learns to factor out nuisance variables

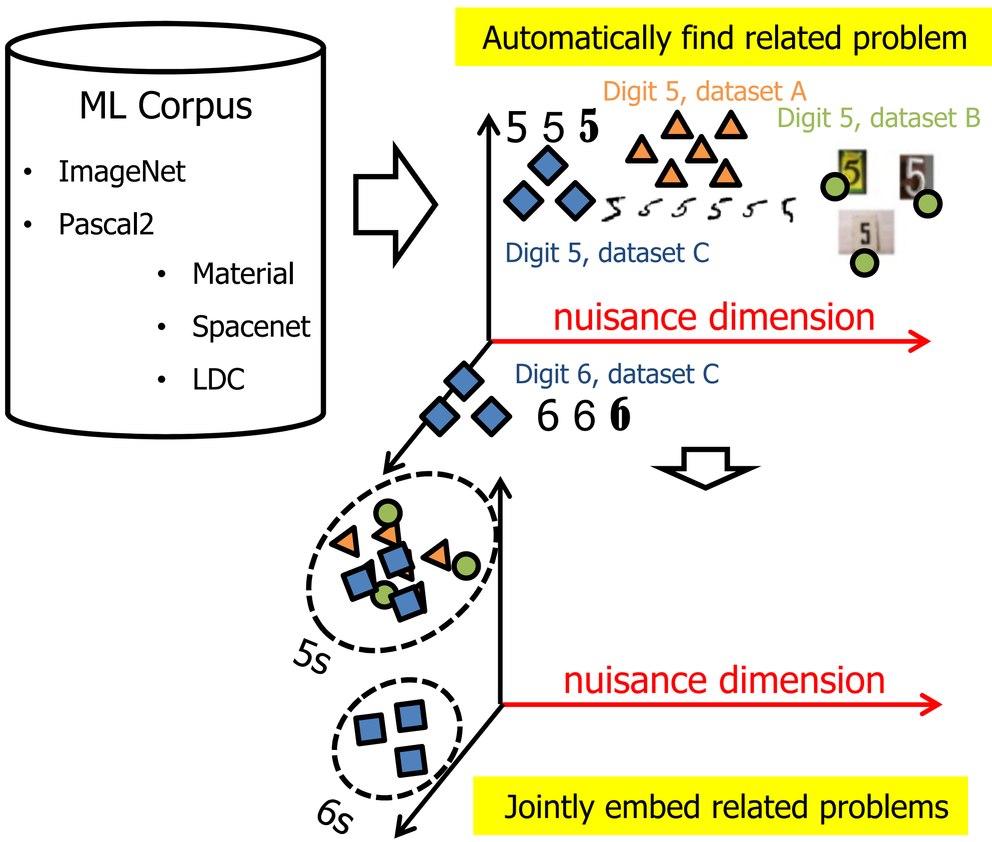

2. Learn what matters via **automated** and optimal experimentation

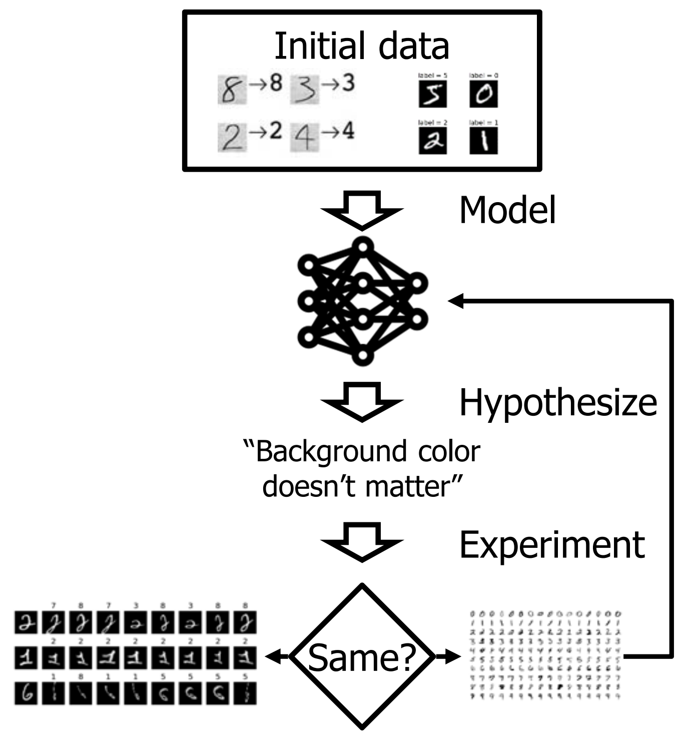

TA1: Automatically derive learning objectives

- **Goal:** Learn correlated objective function with small numbers of examples

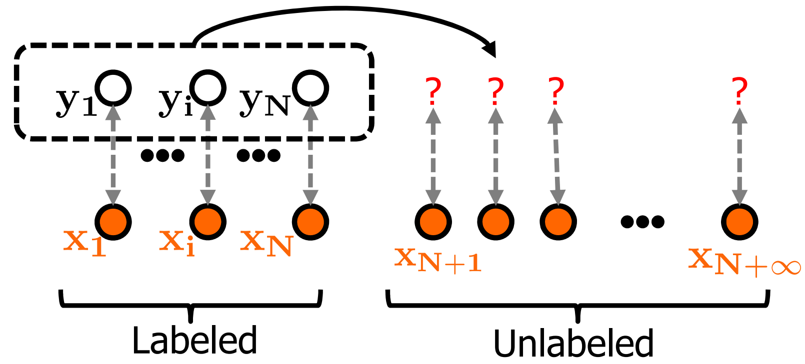

- **Approaches:**

1. **Automated** discovery of correlated reward

Goal: object detection

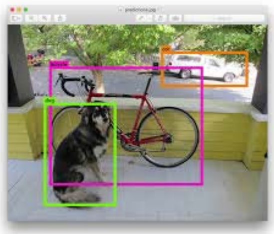

Find correlated reward problem

robot navigation

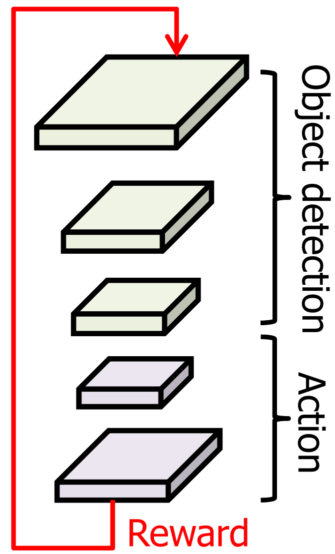

Deep Q Network	81.4%
Object DNN	64.3%

D. Reddy, "LSD-Net: Look, Step and Detect for Joint Navigation and Multi-View Recognition with Deep Reinforcement Learning", ICLR 2018

2. **Automatically discover** and learn low-resource generators to train DNNs

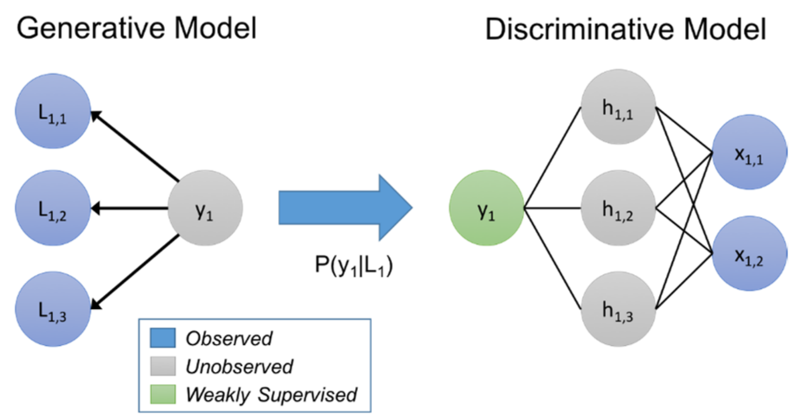

10^4 human + 10^8 generated	72%
10^4 human labels	32%
10^9 human labels	76.4%

A. Ratner et al, "Snorkel: Rapid Training Data Creation with Weak Supervision", VLDB 2018

TA2: Limits of learning and adaptation

- **Goal:** Analytical model of optimal training and adaptation efficiency in terms of labeled data
  1. What is the “true” rank of a classification problem? How complex is an ML problem?
  2. How much labeled data is required to discover a problem’s true rank?
  3. How complex is the data?
  4. What are the discoverability limits for optimal manifolds given sampled data?

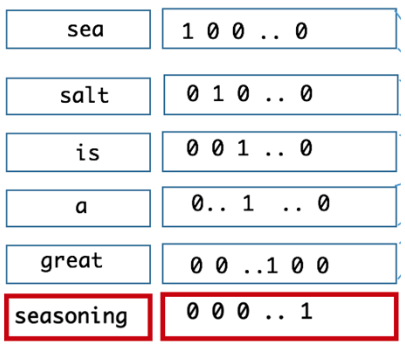

Word prediction (500k dims)

Is this minimal?

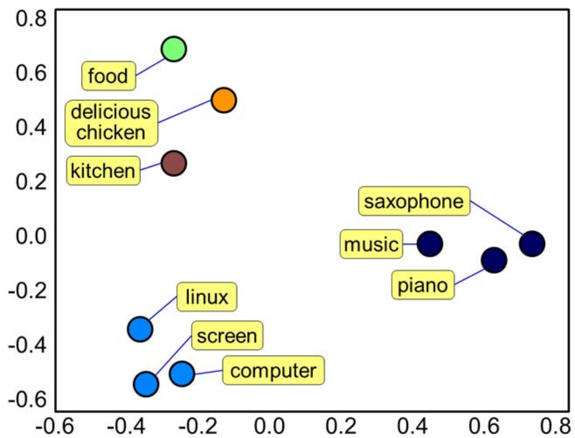

Meaning regression (rank 200)

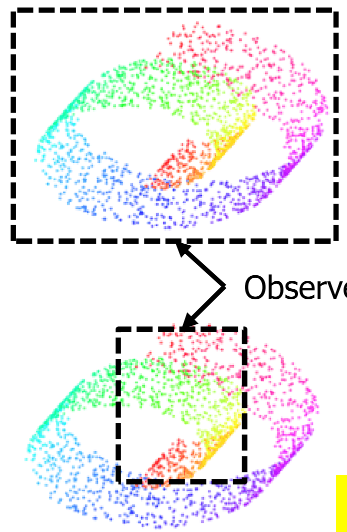

1D manifold

? 1D manifold

Discoverable?

- **Approach:** Formally prove the limits of learning problems
  1. Formally define an approximate problem complexity measure
  2. Formally characterize ML problems and domain mismatch
  3. Prove complexity and transfer limits for problem classes and domains (sample complexity under adaptation, transfer and training)
  4. Prove minimum data rank given M observed samples for chosen accuracy ε (data complexity)

Evaluation and metrics

- Evaluation design
  - 12-month release cycle for challenge problems tied to evaluation
  - Problems: images object detection/recognition, video activity recognition, translation

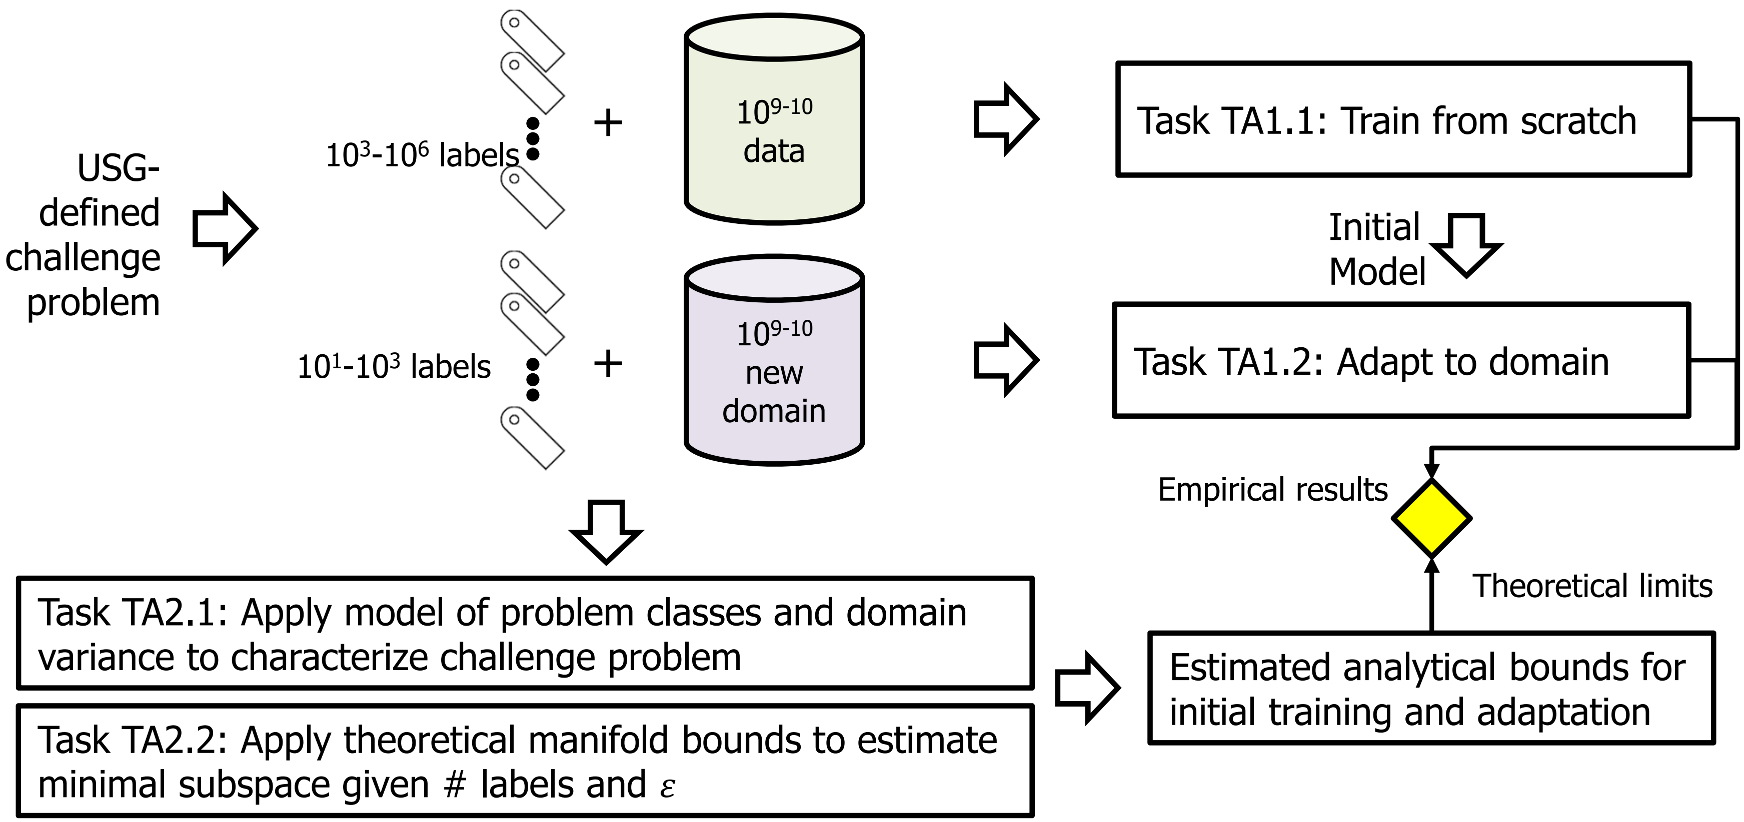

TA1 goals and metrics

LwLL's objective is reduction in labels without loss of state of the art performance; all percentages represent current state of the art

Challenges	Phase 1: Train (TA1.1)	Phase 1: Adapt (TA1.2)	Phase 2: Train (TA1.1)	Phase 2: Adapt (TA1.2)
Object detection[1] Train: LSVRC (open) Adapt: TBD Metric: mAP @ # labels	80% @ $\mathbf{10^5}$	80% @ $\mathbf{10^3}$	80% @ $\mathbf{10^2}$	80% @ $\mathbf{10^2}$
Object classification[2] Train: LSVRC (open) Adapt: TBD Metric: mAP @ # labels	97% @ $\mathbf{10^6}$	97% @ $\mathbf{10^3}$	97% @ $\mathbf{10^3}$	97% @ $\mathbf{10^2}$
Activity recognition[3] Train: TRECVid MED task Adapt: TBD Metric: mAP @ # labels			41% @ $\mathbf{10^3}$	41% @ $\mathbf{10^2}$
Machine translation[4] Train: OpenMT task Adapt: TBD Metric: BLEU @ # labels			47% @ $\mathbf{10^3}$	47% @ $\mathbf{10^2}$

[1] LSVRC 2017 best DET system, unlimited data track
[2] LSVRC 2017 best LOC Task 1b system, unlimited data track
[3] TRECVid 2017 best system MED task
[4] OpenMT 2015 best system Arabic-English task

TA2 goals

Phase 1

Formally define an approximate problem complexity measure
(minimum learning efficiency)

Formally characterize ML problems and domain mismatch

Phase 2

Find and prove the properties of a data sample M essential to satisfying the approximate Johnson-Lindenstrauss transform for training and adaptation

Prove approximate data complexity bounds
(e.g. approximate Johnson-Lindenstrauss $O(\varepsilon^{-2} log N)$ given $M \ll N$)

Prove bounds on problem isomorphism for transfer and adaptation
(transfer limits)

Prove complexity limits for problem classes and domains
(sample complexity)

GOAL: Analytical model of optimal training and adaptation efficiency in terms of labeled data

Use model to estimate bounds

Challenge problem #1

Challenge problem #2

Challenge problem #3

Schedule

Execution and phasing

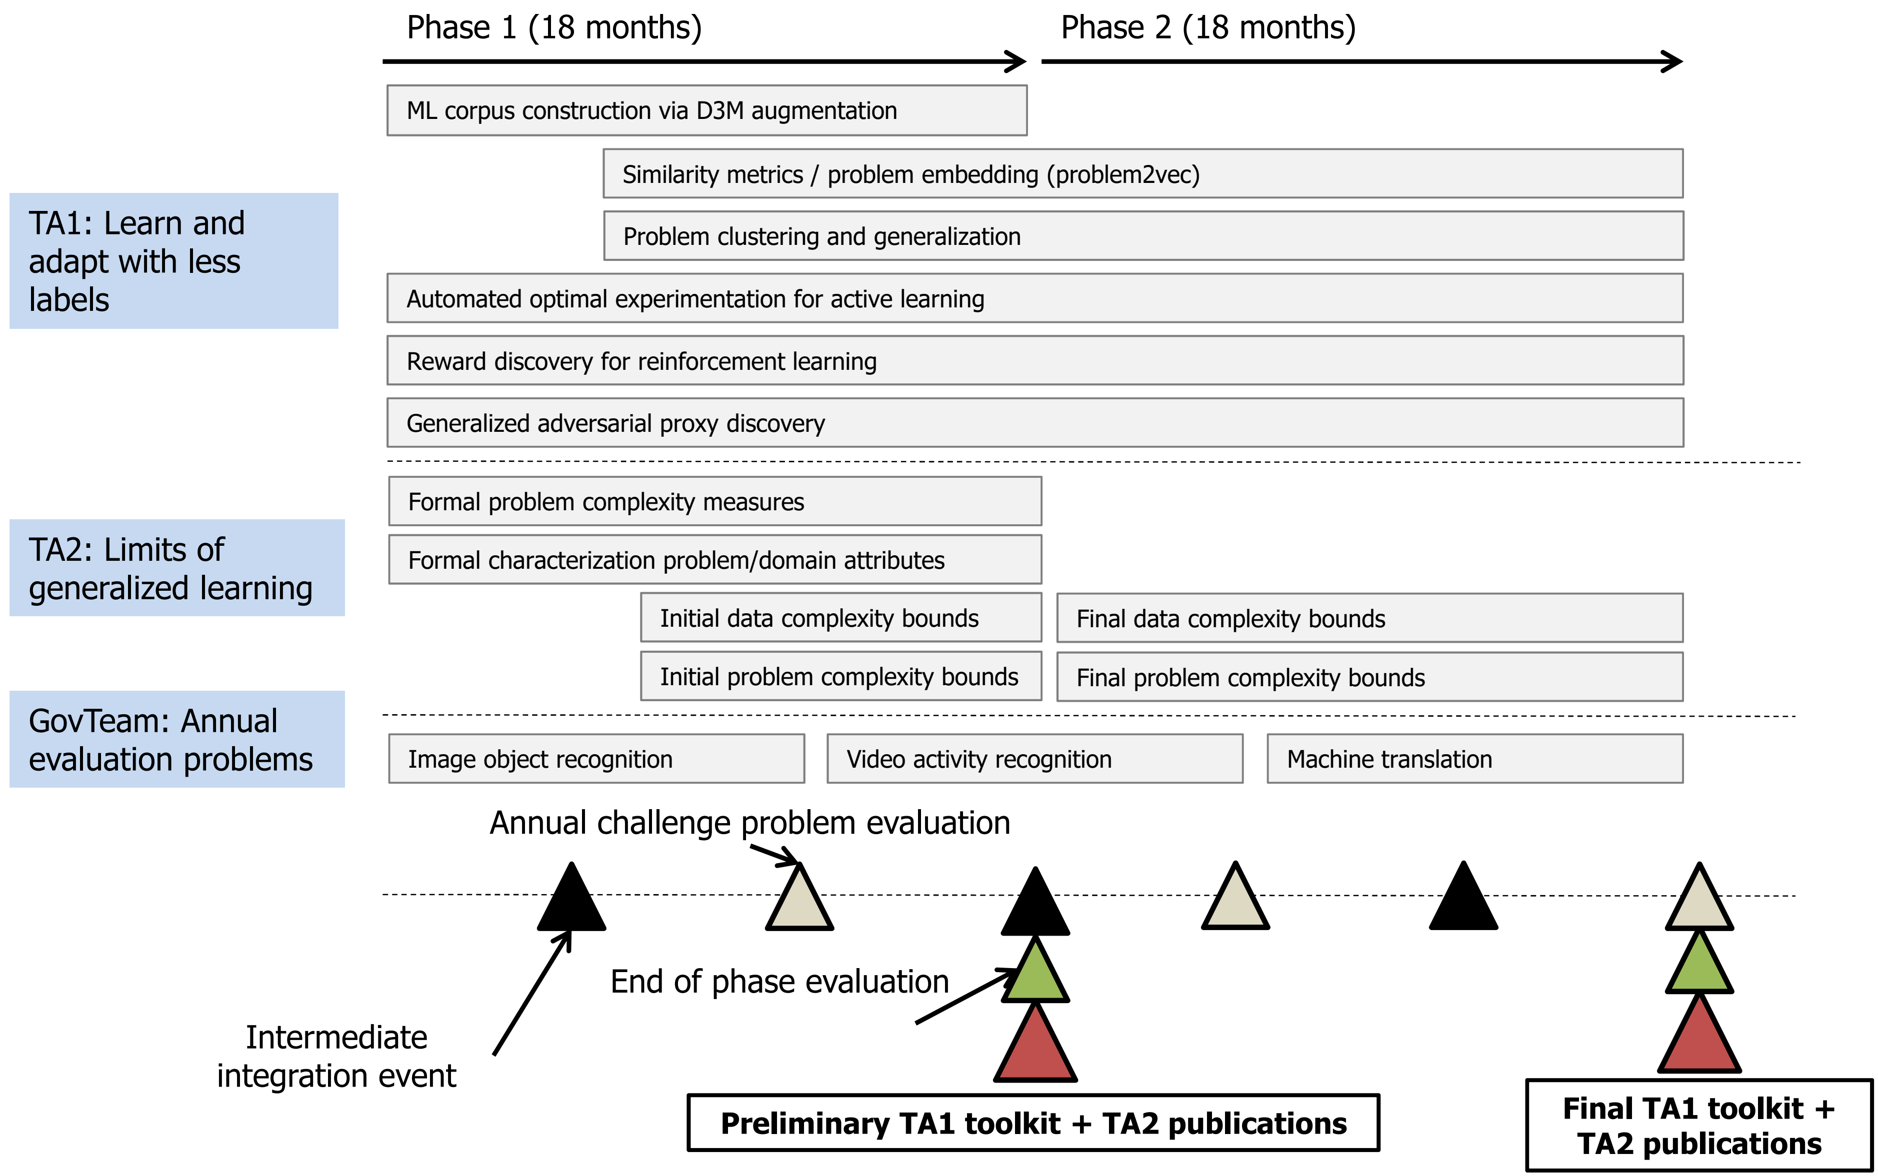

www.darpa.mil