AIFARMS
Artificial Intelligence for Future Agricultural Resilience, Management, and Sustainability

Adaptive Knowledge Transfer on Evolving Domains

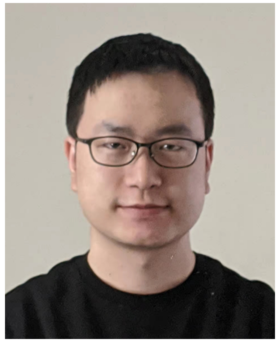

Jun Wu
University of Illinois Urbana-Champaign

Hanghang Tong
University of Illinois Urbana-Champaign

Elizabeth Ainsworth
University of Illinois Urbana-Champaign

Jingrui He
University of Illinois Urbana-Champaign

Presented by **Jun Wu**
Email: junwu3@illinois.edu
Homepage: https://publish.illinois.edu/junwu3/

ILLINOIS

Roadmap

Background
- Dynamic transfer learning
- Assumptions

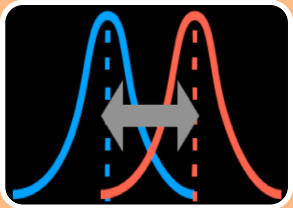

Methodology
- Label-informed $\mathcal{C}$-divergence
- Generalization error bound
- Adversarial VAE framework

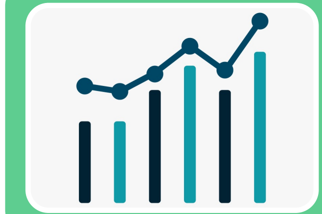

Experiments
- Performance comparison
- Ablation study

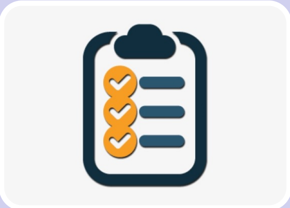

Conclusion
- Algorithm & analysis & evaluation

Background

- Transfer Learning
  - **Input**: A source domain and a target domain
  - **Output**: Prediction function on the target domain
- Assumptions
  - **Relatedness**: Domains are distributionally similar.
  - **Static domains**: All the domain data are static

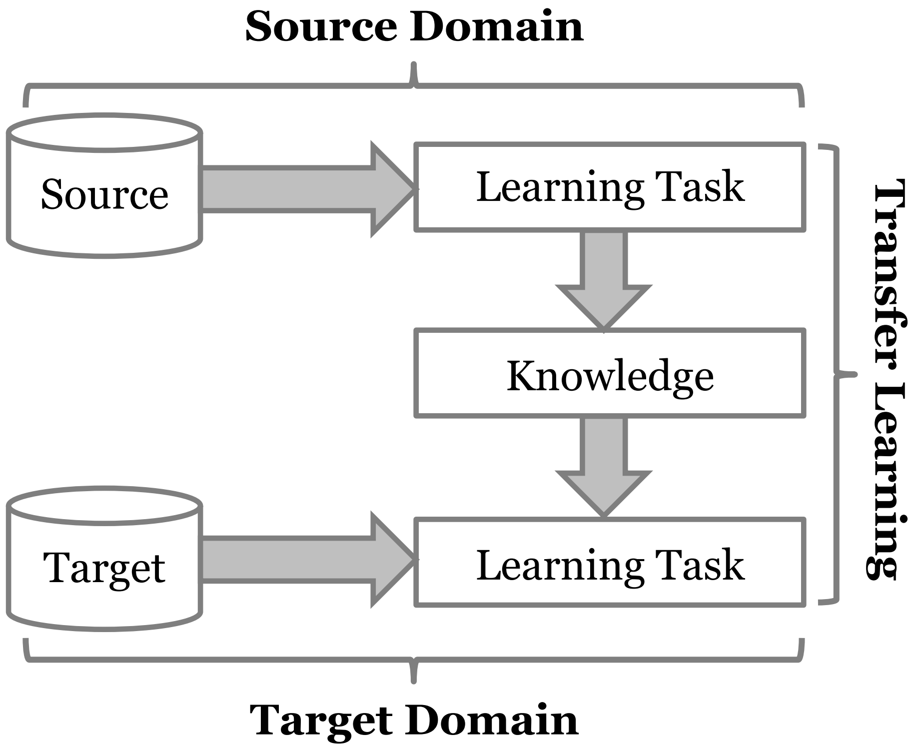

Behnam Neyshabur, et al. "**What is being transferred in transfer learning**?." NeurIPS 2020.

Background

- Transfer Learning
  - **Input**: A source domain and a target domain
  - **Output**: Prediction function on the target domain

- Assumptions
  - **Relatedness**: Domains are distributionally similar.
  - **Static domains**: All the domain data are static

- Applications
  - **Computer vision**: Object recognition
  - **Natural language processing**: Sentiment analysis
  - **Graph mining**: Traffic flow prediction
  - **Agriculture analysis**: Plant phenotyping

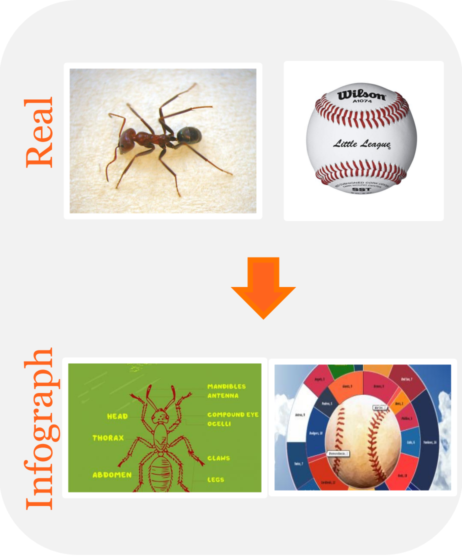

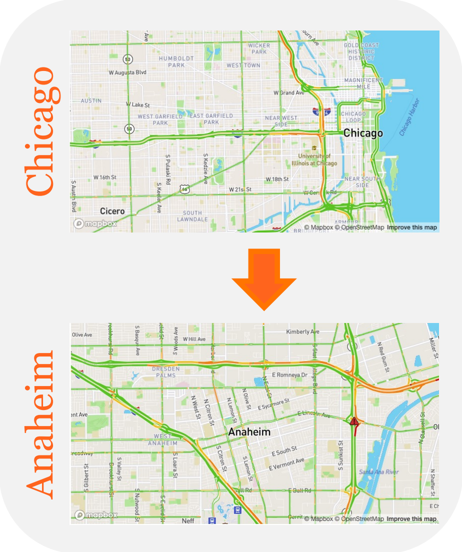

Ozan Sener, et al. "**Learning transferrable representations for unsupervised domain adaptation.**" NeurIPS 2016.
Junteng Jia, et al. "**Residual correlation in graph neural network regression.**" KDD 2020.
Jun Wu et al. "**Adaptive transfer learning for plant phenotyping.**" MLCA 2021

Problem Definition

- Dynamic transfer learning
  - Input: A static source domain (fully labeled);
    A time-evolving target domain (limited labeled and adequate unlabeled)
  - Goal: **Learn a prediction function for the newest target domain**
  - Assumptions
    - **Relatedness**: The source and initial target domains are distributionally related.
    - **Evolvement**: The target domain is continuously evolving over time.

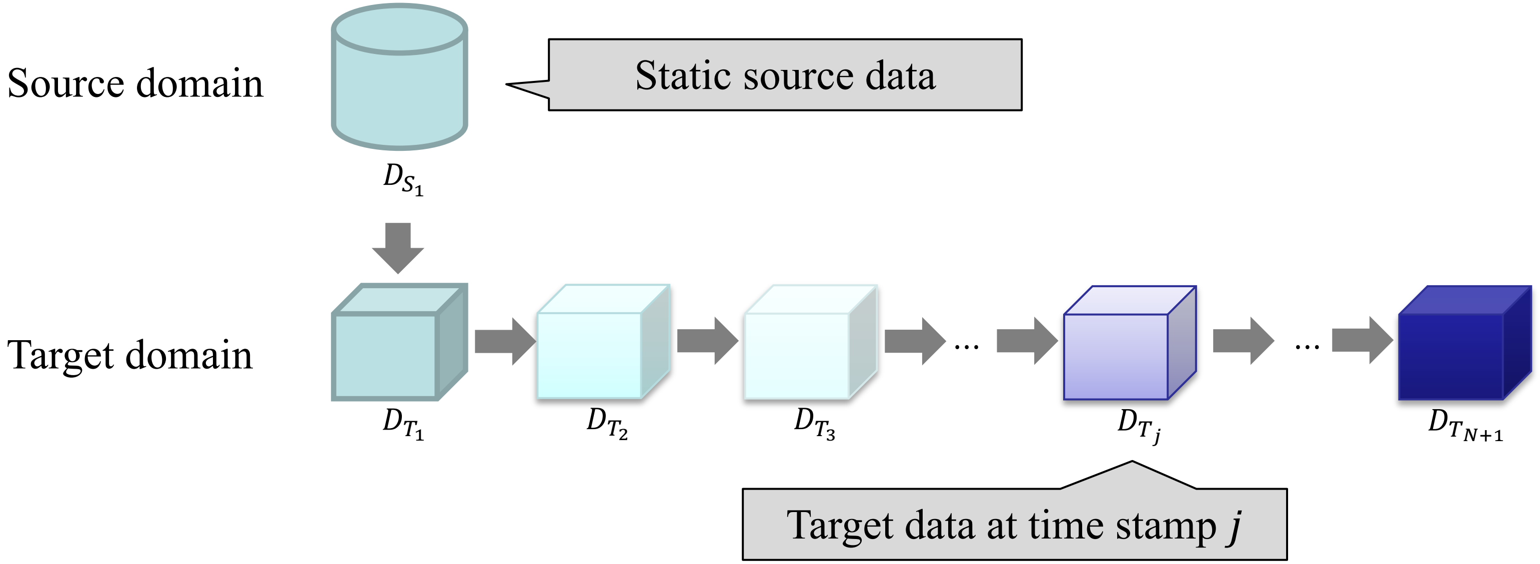

Judy Hoffman, et al. "**Continuous manifold based adaptation for evolving visual domains**." CVPR 2014.
Ananya Kumar, et al. "**Understanding self-training for gradual domain adaptation**." ICML 2020.

Roadmap

Background
- Dynamic transfer learning
- Assumptions

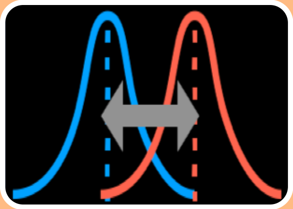

Methodology
- Label-informed $\mathcal{C}$-divergence
- Generalization error bound
- Adversarial VAE framework

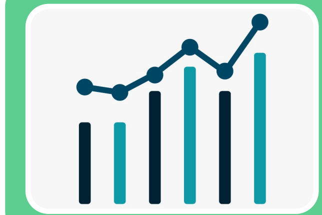

Experiments
- Performance comparison
- Ablation study

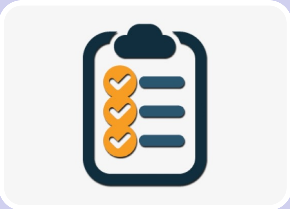

Conclusion
- Algorithm & analysis & evaluation

Label-Informed Domain Discrepancy

- Label-informed domain discrepancy ($\mathcal{C}$-divergence):
  - Key idea: Measure the **joint distribution difference** over both input and output spaces

$$d_{\mathcal{C}}(\mathcal{D}_S, \mathcal{D}_T) = \sup_{h \in \mathcal{H}} \left| \Pr_{\mathcal{D}_S}[\{I(h), y=1\} \cup \{\overline{I(h)}, y=0\}] - \Pr_{\mathcal{D}_T}[\{I(h), y=1\} \cup \{\overline{I(h)}, y=0\}] \right|$$

- Comparison with existing marginal discrepancy measures, e.g., $\mathcal{A}$-distance [1]

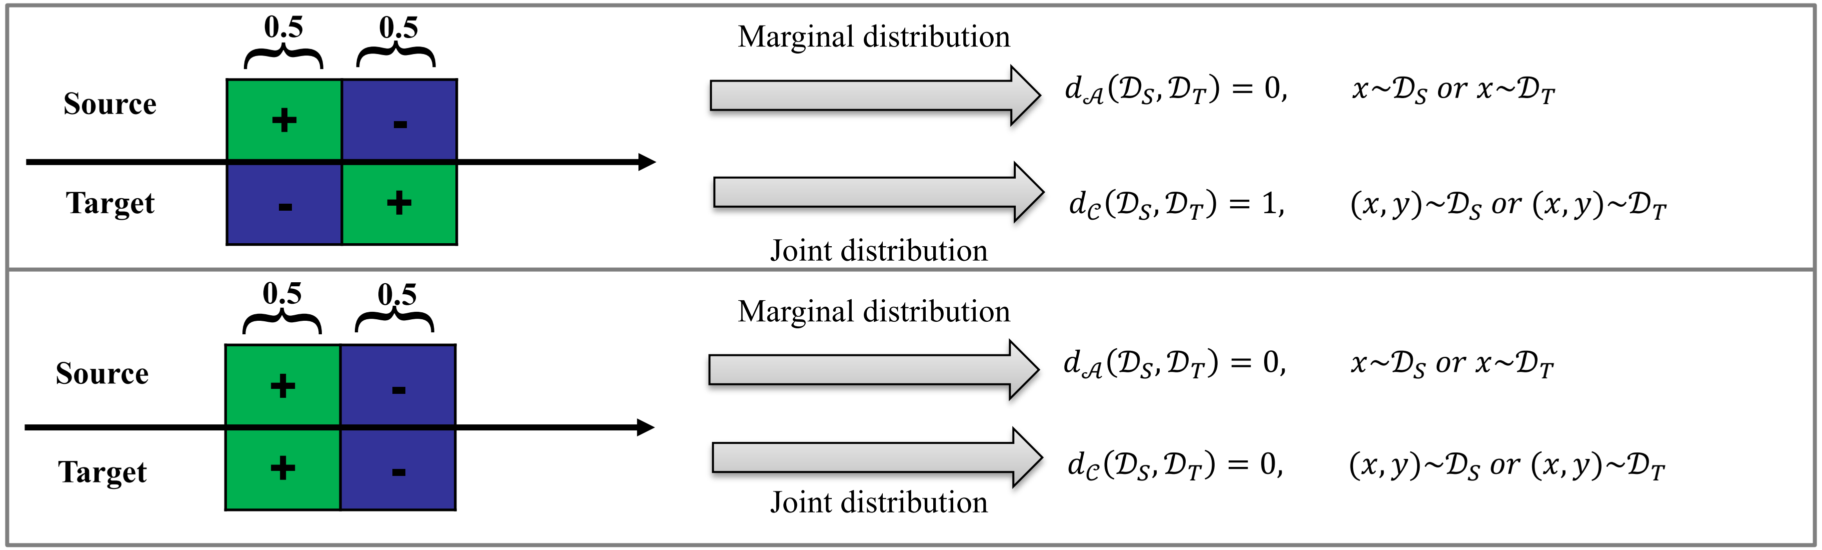

[1] Shai Ben-David, et al. "**A theory of learning from different domains**." *Machine learning* 2010.

Theoretical Analysis

- The expected target error $\epsilon_{T_{t+1}}$ on time stamp t+1 is bounded by
  - (i) Empirical classification error $\hat{\epsilon}_{T_0}$ on source data
  - (ii) Empirical classification error $\hat{\epsilon}_{T_j}$ $(j = 1, \ldots, t)$ on historical target data
  - (iii) Label-informed domain discrepancy $d_{\mathcal{C}}$ among domains

Theorem: Assume the loss function $L(\cdot,\cdot)$ is bounded. Given a source domain $\mathcal{D}_S$ (denoted as $\mathcal{D}_{T_0}$) and historical target domain $\left\{\mathcal{D}_{T_j}\right\}_{j=1}^{t}$, for $h \in \mathcal{H}$ and $\delta \in (0,1)$, with probability at least $1-\delta$, the target domain error $\epsilon_{T_{t+1}}$ on the newest target domain is bounded as follows.

$$\epsilon_{T_{t+1}}(h) \leq \frac{1}{\bar{\mu}}\left(\sum_{j=0}^{t} \mu^{t-j}\hat{\epsilon}_{T_j}(h) + M\sum_{j=0}^{t}\mu^{t-j} d_{\mathcal{C}}\left(\widehat{\mathcal{D}}_{T_j}, \widehat{\mathcal{D}}_{T_{t+1}}\right) + M\Lambda\right)$$

where Λ is a Rademacher complexity term.

Proposed Algorithm: TransLATE

- TransLATE: **Trans**fer learning with **la**bel-informed dis**t**ribution alignm**e**nt

- Objective function

$$\mathcal{J} = \sum_{j=0}^{t} \mu^{t-j} \Big(\underbrace{L_{clc}(T_j, T_{t+1})}_{\text{Empirical error}} + \overbrace{d_{\mathcal{C}}\left(\widehat{\mathcal{D}}_{T_j}, \widehat{\mathcal{D}}_{T_{t+1}}\right)}^{\text{Label-informed domain discrepancy}} + \underbrace{\lambda L_{ELBO}(T_j, T_{t+1})}_{\text{Variational Autoencoder for feature extraction}} \Big)$$

  - Assume data distribution is **continuously evolving**, i.e., $0 \leq \mu \leq 1$
  - **Variational autoencoder** learns the common feature space
  - The empirical error on historical target task is iteratively estimated

Roadmap

Background
- Dynamic transfer learning
- Assumptions

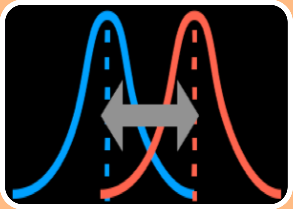

Methodology
- Label-informed $\mathcal{C}$-divergence
- Generalization error bound
- Adversarial VAE framework

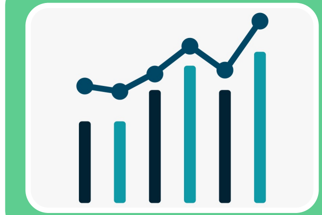

Experiments
- Performance comparison
- Ablation study

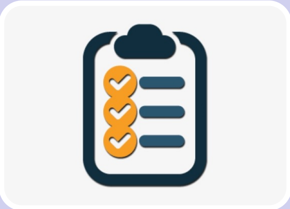

Conclusion
- Algorithm & analysis & evaluation

Experiments

- Data sets
  - Synthetic
  - Office-31
  - Office-Home
- Metric
  - Classification accuracy
- Baselines
  - Single domain: SourceOnly, TargetERM
  - Static adaptation: DAN, DANN, and MDD
  - Dynamic adaptation: CUA, GST, and CIDA

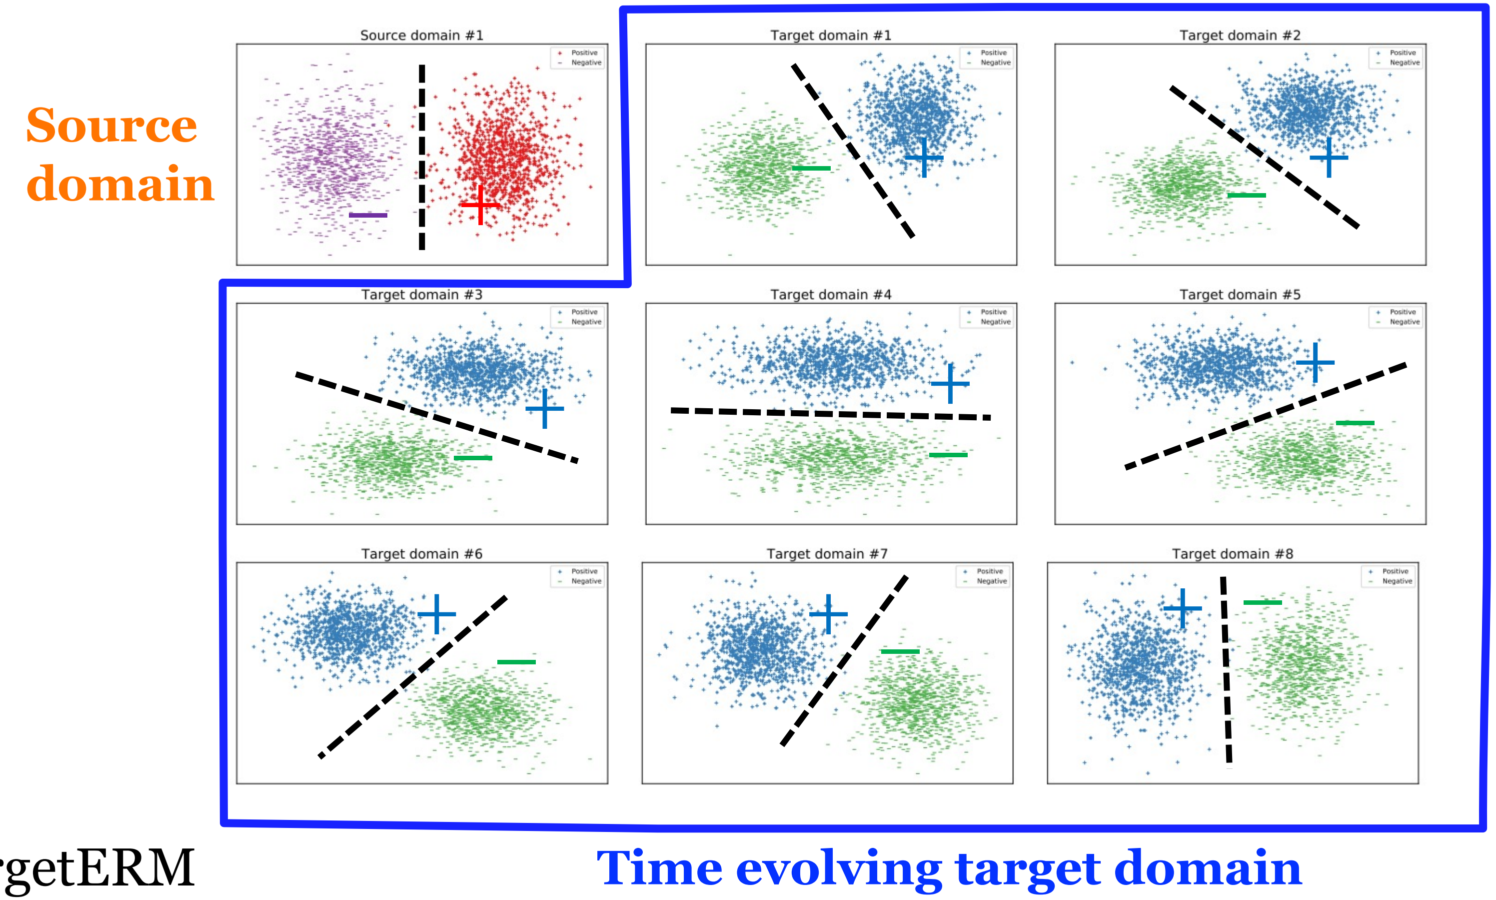

Ananya Kumar, et al. "**Understanding self-training for gradual domain adaptation**." ICML 2020.
Hao Wang, et al. "**Continuously indexed domain adaptation**." ICML 2020.
Yuchen Zhang, et al. "**Bridging theory and algorithm for domain adaptation**." ICML 2019.

Model Analysis

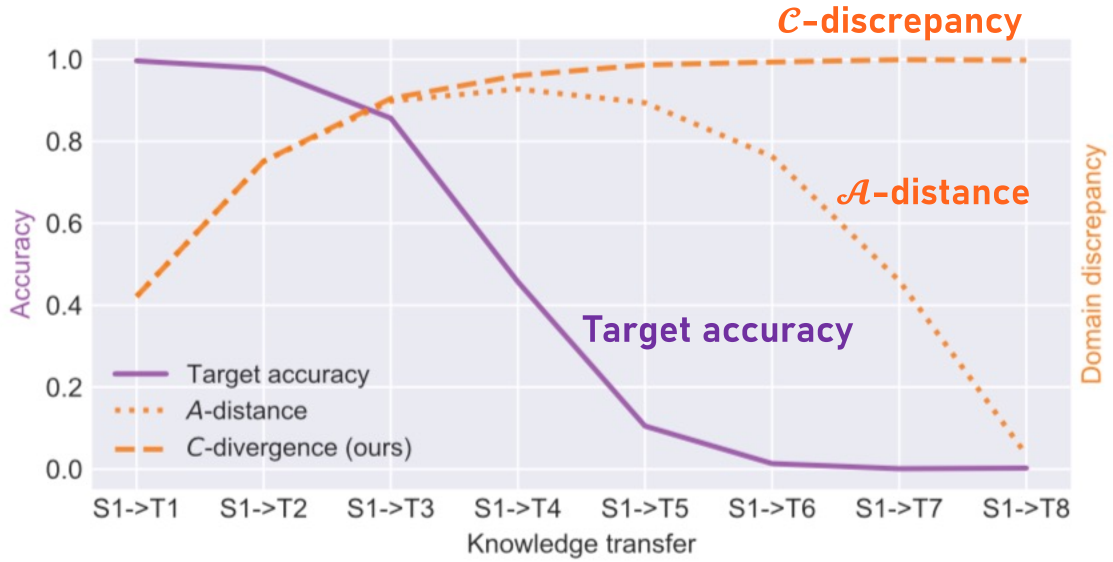

Comparison of domain discrepancy and target error

Comparison of different error bounds

Observations:

(a) Compared to $\mathcal{A}$-distance [1], $\mathcal{C}$-divergence better characterizes the transferability from the source to the target domains

(b) Our $\mathcal{C}$-divergence based error bound is much tighter than the baseline [1] based on $\mathcal{A}$-distance

[1] Shai Ben-David, et al. "**A theory of learning from different domains**." *Machine learning* 2010.

Results

Observation: TransLATE achieves significantly better performance (+10%) on the newest target domain.

Classification accuracy

Office-31

	Amazon → Webcam					DSLR → Webcam				
	T1	T2	T3	T4	T5	T1	T2	T3	T4	T5
SourceOnly	0.7490	0.2255	0.2282	0.1275	0.1503	0.9651	0.4309	0.3329	0.1611	0.2027
TargetERM	0.5584	0.3933	0.4215	0.3396	0.3732	0.4966	0.4201	0.4188	0.3248	0.4067
DAN [9]	0.8537	0.5007	0.4993	0.3638	0.4470	0.9772	0.7302	0.6161	0.4765	0.5302
DANN [3]	0.8389	0.4993	0.4121	0.3973	0.3382	0.9651	0.7356	0.6416	0.4510	0.5490
MDD [19]	0.8940	0.6738	0.5490	0.5141	0.4295	0.9724	0.8738	0.7315	0.5047	0.5289
TransLATE$_\infty$	**0.9154**	0.6376	0.5758	0.4591	0.4846	0.9785	0.8591	0.7289	0.4926	0.5557
CUA [2]	0.8349	0.6805	0.6389	0.6456	0.6805	**0.9852**	0.8805	0.8792	0.8362	0.8617
GST [6]	0.8456	0.5987	0.6013	0.5584	0.5960	0.9739	0.8376	0.8134	0.7570	0.7865
CIDA [14]	0.8805	0.7638	0.7624	0.7195	0.7476	0.9812	0.8577	0.8376	0.7973	0.7960
TransLATE	**0.9154**	**0.8134**	**0.8081**	**0.7611**	**0.7826**	0.9785	**0.9235**	**0.9208**	**0.8886**	**0.9154**

Office-Home

	Art → Real World					Clipart → Product				
	T1	T2	T3	T4	T5	T1	T2	T3	T4	T5
SourceOnly	0.7220	0.3947	0.3135	0.2650	0.3512	0.5944	0.1866	0.1342	0.1074	0.1550
TargetERM	0.5643	0.3297	0.3010	0.2582	0.3299	0.6033	0.3805	0.4232	0.3061	0.3791
DAN [9]	0.7341	0.4901	0.4193	0.3686	0.4597	0.7186	0.4201	0.3921	0.3352	0.4113
DANN [3]	0.7359	0.5092	0.4155	0.3850	0.4686	0.7063	0.4440	0.3694	0.3343	0.4303
MDD [19]	0.7435	0.5056	0.4331	0.3874	0.4686	0.7264	0.4765	0.3886	0.3514	0.4294
TransLATE$_\infty$	**0.7560**	0.5273	0.4575	0.4080	0.4850	**0.7411**	0.5017	0.4436	0.3634	0.4595
CUA [2]	0.7370	0.5732	0.5181	0.4932	0.5372	0.7143	0.4922	0.4431	0.4310	0.4879
GST [6]	0.7367	0.5283	0.4795	0.4681	0.4826	0.7285	0.5232	0.4782	0.4531	0.4943
CIDA [14]	0.7420	0.5643	0.4983	0.4896	0.5130	0.7226	0.5076	0.4334	0.4030	0.4362
TransLATE	**0.7560**	**0.6046**	**0.5447**	**0.5097**	**0.5459**	**0.7411**	**0.5747**	**0.5318**	**0.5009**	**0.5422**

Roadmap

Background
- Dynamic transfer learning
- Assumptions

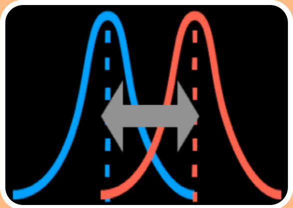

Methodology
- Label-informed $\mathcal{C}$-divergence
- Generalization error bound
- Adversarial VAE framework

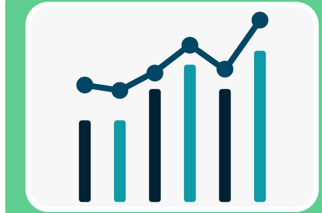

Experiments
- Performance comparison
- Ablation study

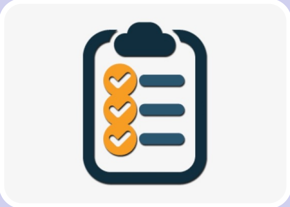

Conclusion
- Algorithm & analysis & evaluation

Conclusion

- **Problem**: Dynamic transfer learning with time evolving target domain
- **Analysis**: Generalization error bounds with the proposed $\mathcal{C}$-divergence
- **Algorithm**: Adversarial variational autoencoder framework based on empirical $\mathcal{C}$-divergence
- **Evaluation**: Competitive performance on modeling the newest target domain

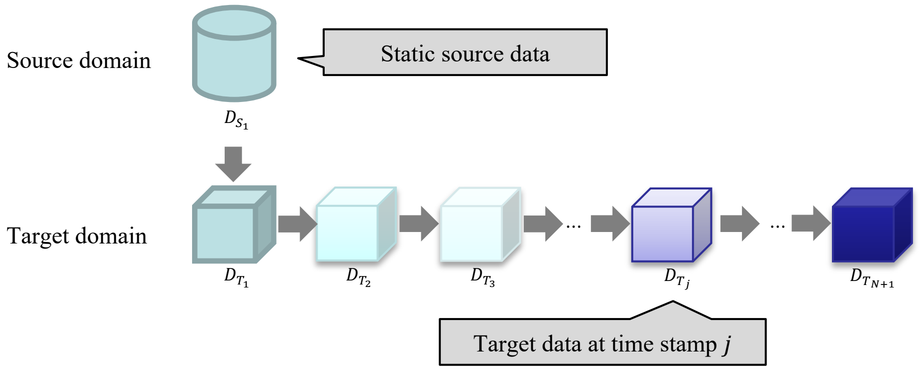

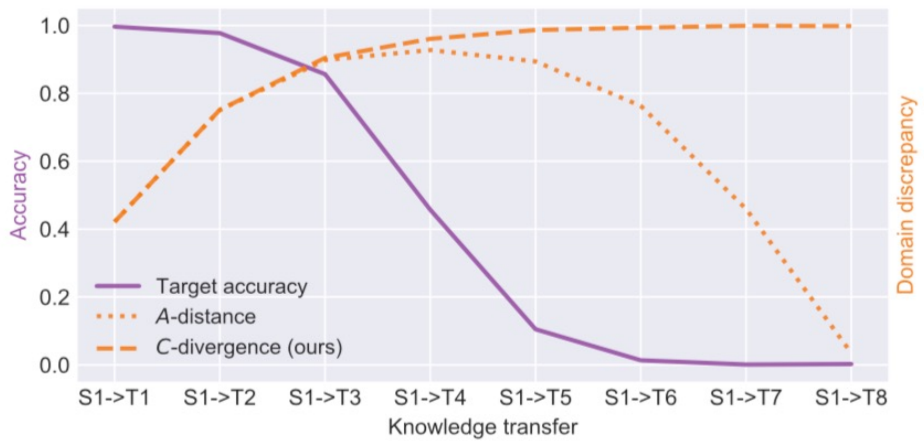

NSF

AIFARMS
Artificial Intelligence for Future Agricultural Resilience, Management, and Sustainability

Adaptive Knowledge Transfer on Evolving Domains

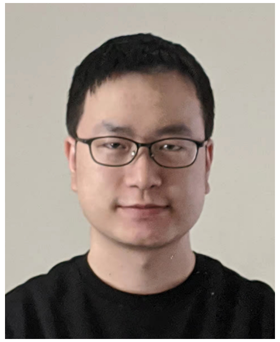

Jun Wu

University of Illinois Urbana-Champaign

Hanghang Tong

University of Illinois Urbana-Champaign

Elizabeth Ainsworth

University of Illinois Urbana-Champaign

Jingrui He

University of Illinois Urbana-Champaign

Please feel free to contact me if you have any question.

Email: junwu3@illinois.edu

Homepage: https://publish.illinois.edu/junwu3/

ILLINOIS