MAPPING CONDITIONAL DISTRIBUTIONS FOR DOMAIN ADAPTATION UNDER GENERALIZED TARGET SHIFT

ICLR 2022

Monday 25th April, 2022 to Friday 29th April, 2022

Matthieu Kirchmeyer[1,2], Alain Rakotomamonjy[2,3], Emmanuel de Bézenac[1], Patrick Gallinari[1,2]

[1]Sorbonne Université, [2]Criteo AI Lab, [3]Université de Rouen

Introduction

Problem definition

SCIENCES SORBONNE UNIVERSITÉ

Unsupervised Domain Adaptation (UDA)

Train a good classifier $h : \mathcal{X} \to \mathcal{Y}$ for an **unlabelled** target domain T given a labelled source domain S.

Problem definition

Unsupervised Domain Adaptation (UDA)

Train a good classifier $h : \mathcal{X} \rightarrow \mathcal{Y}$ for an **unlabelled** target domain T given a labelled source domain S.

$$p_S(X, Y) \neq p_T(X, Y)$$

Problem definition

SCIENCES SORBONNE UNIVERSITÉ

Unsupervised Domain Adaptation (UDA)

Train a good classifier $h : \mathcal{X} \to \mathcal{Y}$ for an **unlabelled** target domain T given a labelled source domain S.

$$p_S(X, Y) \neq p_T(X, Y)$$

Generalized Target Shift

We consider the most challenging UDA setting, Generalized target shift (`GeTarS`) Zhang et al. 2013 where:

$$p_S(X|Y) \neq p_T(X|Y), p_S(Y) \neq p_T(Y)$$

UDA under generalized target shift

SCIENCES SORBONNE UNIVERSITÉ

Standard is to learn domain-invariant representations

$h : \mathcal{X} \overset{g}{\mapsto} \mathcal{Z} \overset{f_S}{\mapsto} \mathcal{Y}$ where g: encoder, f_S: classifier,

UDA under generalized target shift

Standard is to learn domain-invariant representations

$h : \mathcal{X} \overset{g}{\mapsto} \mathcal{Z} \overset{f_S}{\mapsto} \mathcal{Y}$ where g: encoder, f_S: classifier,

Solve two steps jointly:

UDA under generalized target shift

Standard is to learn domain-invariant representations

$h : \mathcal{X} \overset{g}{\mapsto} \mathcal{Z} \overset{f_S}{\mapsto} \mathcal{Y}$ where g: encoder, f_S: classifier,

Solve two steps jointly:

- Learn f_S to classify encoded S samples.

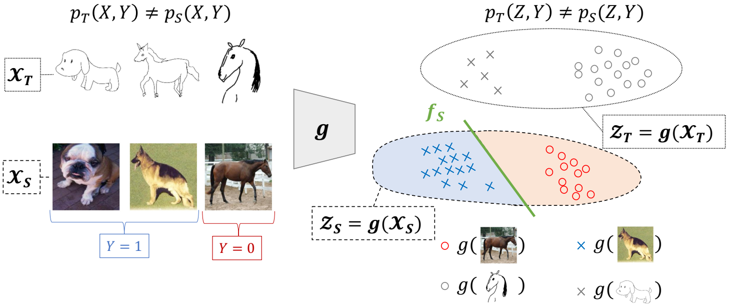

UDA under generalized target shift

SCIENCES SORBONNE UNIVERSITÉ

Standard is to learn domain-invariant representations

$h : \mathcal{X} \overset{g}{\mapsto} \mathcal{Z} \overset{f_S}{\mapsto} \mathcal{Y}$ where g: encoder, f_S: classifier,

Solve two steps jointly:

- Learn f_S to classify encoded S samples.
- Learn g to match encoded S and T samples.

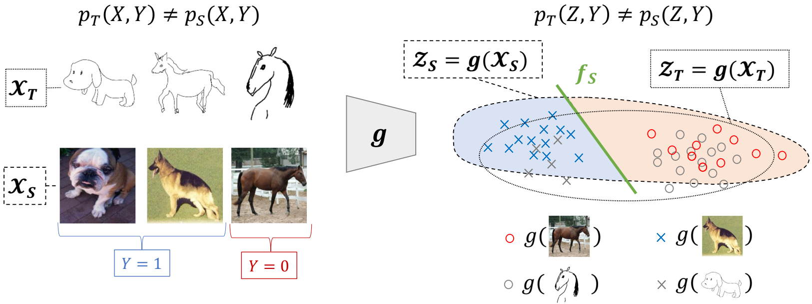

UDA under generalized target shift

Standard is to learn domain-invariant representations

$h : \mathcal{X} \overset{g}{\mapsto} \mathcal{Z} \overset{f_S}{\mapsto} \mathcal{Y}$ where g: encoder, f_S: classifier,

Solve two steps jointly:

- Learn f_S to classify encoded S samples.
- Learn g to match encoded S and T samples.
- Under `GeTarS`, reweight S samples by estimated class-ratios.

Combes et al. 2020; Gong et al. 2016; Rakotomamonjy et al. 2021; Shui et al. 2021

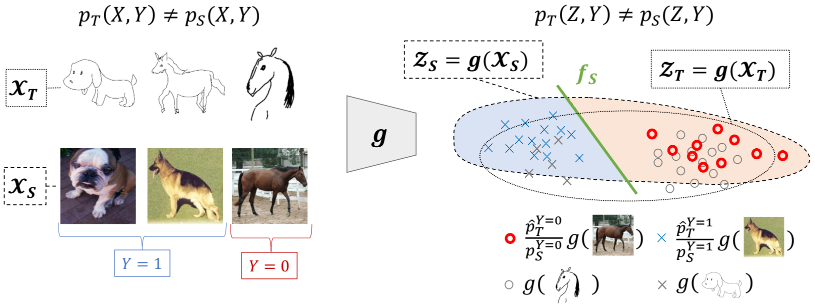

UDA under generalized target shift

Limitations of invariant representation learning

- It is prone to instabilities due to adversarial alignment, especially without established NN architectures.

UDA under generalized target shift

SCIENCES SORBONNE UNIVERSITÉ

Limitations of invariant representation learning

- It is prone to instabilities due to adversarial alignment, especially without established NN architectures.
- Target discriminativity may degrade Liu et al. 2019.

UDA under generalized target shift

Limitations of invariant representation learning

- It is prone to instabilities due to adversarial alignment, especially without established NN architectures.
- Target discriminativity may degrade Liu et al. 2019.
- Generalization guarantees for `GeTarS` are derived under strong assumptions which may not hold in practice.

UDA under generalized target shift

SCIENCES SORBONNE UNIVERSITÉ

Limitations of invariant representation learning

- It is prone to instabilities due to adversarial alignment, especially without established NN architectures.
- Target discriminativity may degrade Liu et al. 2019.
- Generalization guarantees for `GeTarS` are derived under strong assumptions which may not hold in practice.

OSTAR: an alternative to domain-invariance for `GeTarS`

UDA under generalized target shift

SCIENCES SORBONNE UNIVERSITÉ

Limitations of invariant representation learning

- It is prone to instabilities due to adversarial alignment, especially without established NN architectures.
- Target discriminativity may degrade Liu et al. 2019.
- Generalization guarantees for `GeTarS` are derived under strong assumptions which may not hold in practice.

OSTAR: an alternative to domain-invariance for `GeTarS`

- Aligns pretrained representations with a NN mapping.

UDA under generalized target shift

SCIENCES SORBONNE UNIVERSITÉ

Limitations of invariant representation learning

- It is prone to instabilities due to adversarial alignment, especially without established NN architectures.
- Target discriminativity may degrade Liu et al. 2019.
- Generalization guarantees for `GeTarS` are derived under strong assumptions which may not hold in practice.

OSTAR: an alternative to domain-invariance for `GeTarS`

- Aligns pretrained representations with a NN mapping.
- This mapping is regularized with Optimal Transport (OT):

UDA under generalized target shift

SCIENCES SORBONNE UNIVERSITÉ

Limitations of invariant representation learning

- It is prone to instabilities due to adversarial alignment, especially without established NN architectures.
- Target discriminativity may degrade Liu et al. 2019.
- Generalization guarantees for `GeTarS` are derived under strong assumptions which may not hold in practice.

OSTAR: an alternative to domain-invariance for `GeTarS`

- Aligns pretrained representations with a NN mapping.
- This mapping is regularized with Optimal Transport (OT):
  - useful inductive biases for stability and performance.

UDA under generalized target shift

Limitations of invariant representation learning

- It is prone to instabilities due to adversarial alignment, especially without established NN architectures.
- Target discriminativity may degrade Liu et al. 2019.
- Generalization guarantees for `GeTarS` are derived under strong assumptions which may not hold in practice.

OSTAR: an alternative to domain-invariance for `GeTarS`

- Aligns pretrained representations with a NN mapping.
- This mapping is regularized with Optimal Transport (OT):
  - useful inductive biases for stability and performance.
  - under mild assumptions, provides two theoretical guarantees:

UDA under generalized target shift

Limitations of invariant representation learning

- It is prone to instabilities due to adversarial alignment, especially without established NN architectures.
- Target discriminativity may degrade Liu et al. 2019.
- Generalization guarantees for `GeTarS` are derived under strong assumptions which may not hold in practice.

OSTAR: an alternative to domain-invariance for `GeTarS`

- Aligns pretrained representations with a NN mapping.
- This mapping is regularized with Optimal Transport (OT):
  - useful inductive biases for stability and performance.
  - under mild assumptions, provides two theoretical guarantees:
    - explicit control of the target risk.

UDA under generalized target shift

SCIENCES SORBONNE UNIVERSITÉ

Limitations of invariant representation learning

- It is prone to instabilities due to adversarial alignment, especially without established NN architectures.
- Target discriminativity may degrade Liu et al. 2019.
- Generalization guarantees for `GeTarS` are derived under strong assumptions which may not hold in practice.

OSTAR: an alternative to domain-invariance for `GeTarS`

- Aligns pretrained representations with a NN mapping.
- This mapping is regularized with Optimal Transport (OT):
  - useful inductive biases for stability and performance.
  - under mild assumptions, provides two theoretical guarantees:
    - explicit control of the target risk.
    - unicity of the solution.

OSTAR framework

Optimal Sample Transport and Reweight (OSTAR) (I)

SCIENCES SORBONNE UNIVERSITÉ

Objectives

- Encode input S and T samples with g

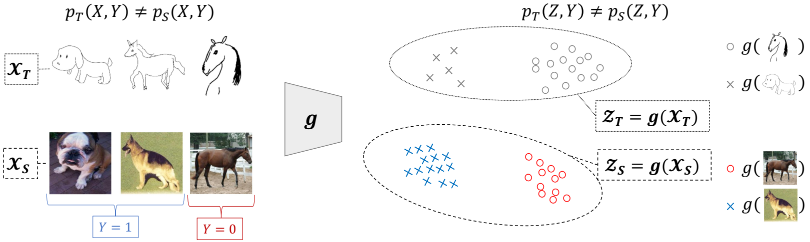

Optimal Sample Transport and Reweight (OSTAR) (I)

Objectives

- With fixed representations jointly
  1. Map encoded S samples onto T with ϕ under OT constraints.

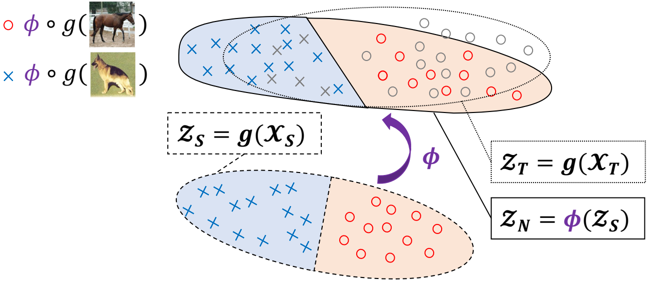

Optimal Sample Transport and Reweight (OSTAR) (I)

SCIENCES SORBONNE UNIVERSITÉ

Objectives

- With fixed representations jointly

2. Reweight mapped S samples with class-ratio estimates $\boldsymbol{p}_N^Y / \boldsymbol{p}_S^Y$.

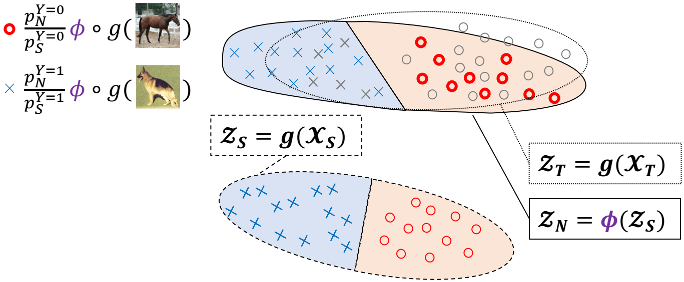

Optimal Sample Transport and Reweight (OSTAR) (I)

SCIENCES SORBONNE UNIVERSITÉ

Objectives

- With fixed representations jointly

3. Train classifier f_N on reweighted and mapped S samples

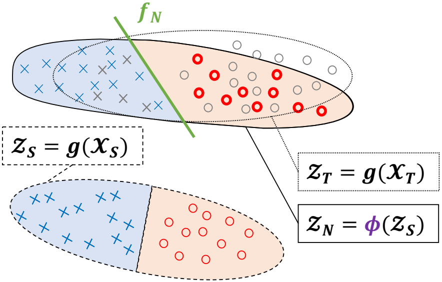

Optimal Sample Transport and Reweight (OSTAR) (I)

Objectives

- Use f_N for inference on encoded T samples.

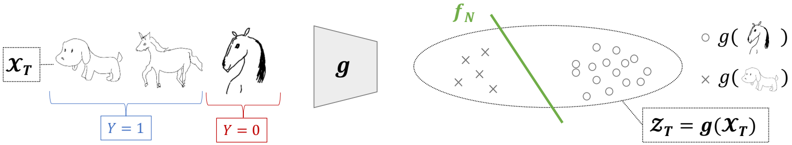

Optimal Sample Transport and Reweight (OSTAR) (II)

SCIENCES SORBONNE UNIVERSITÉ

Improving target discriminativity

- The encoder g was so far fixed.

Optimal Sample Transport and Reweight (OSTAR) (II)

SCIENCES SORBONNE UNIVERSITÉ

Improving target discriminativity

- The encoder g was so far fixed.
- We update g using target pseudo-labels with Information Maximization (IM).

Optimal Sample Transport and Reweight (OSTAR) (II)

SCIENCES SORBONNE UNIVERSITÉ

Improving target discriminativity

- The encoder g was so far fixed.
- We update g using target pseudo-labels with Information Maximization (IM).
- Improves target discriminativity and helps alignment.

Optimal Sample Transport and Reweight (OSTAR) (II)

SCIENCES SORBONNE UNIVERSITÉ

Improving target discriminativity

- The encoder g was so far fixed.
- We update g using target pseudo-labels with Information Maximization (IM).
- Improves target discriminativity and helps alignment.

Theoretical results

- New upper-bound to the target risk under `GeTarS`.

Optimal Sample Transport and Reweight (OSTAR) (II)

Improving target discriminativity

- The encoder g was so far fixed.
- We update g using target pseudo-labels with Information Maximization (IM).
- Improves target discriminativity and helps alignment.

Theoretical results

- New upper-bound to the target risk under `GeTarS`.
- We define mild assumptions on g, for which

Optimal Sample Transport and Reweight (OSTAR) (II)

SCIENCES SORBONNE UNIVERSITÉ

Improving target discriminativity

- The encoder g was so far fixed.
- We update g using target pseudo-labels with Information Maximization (IM).
- Improves target discriminativity and helps alignment.

Theoretical results

- New upper-bound to the target risk under `GeTarS`.
- We define mild assumptions on g, for which
  - `OSTAR` controls the target risk at optimum.

Optimal Sample Transport and Reweight (OSTAR) (II)

SCIENCES SORBONNE UNIVERSITÉ

Improving target discriminativity

- The encoder g was so far fixed.
- We update g using target pseudo-labels with Information Maximization (IM).
- Improves target discriminativity and helps alignment.

Theoretical results

- New upper-bound to the target risk under `GeTarS`.
- We define mild assumptions on g, for which
  - `OSTAR` controls the target risk at optimum.
  - Its solution $(\phi, \boldsymbol{p}_N^Y)$ is unique.

Results

t-SNE feature visualizations

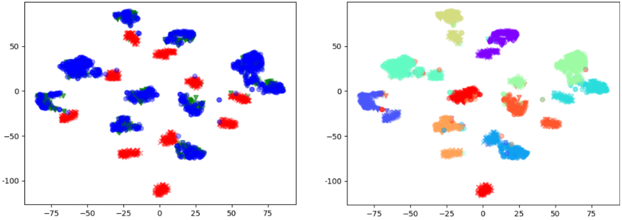

(a) USPS→MNIST

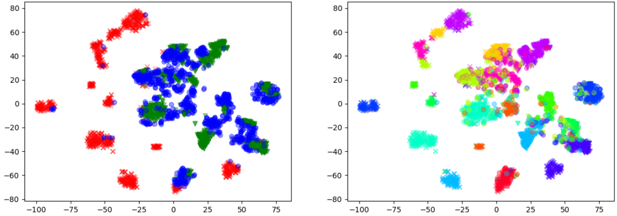

(b) VisDA

(left) source, target and mapped source. (right) classes in source and target.

Performance on visual UDA datasets

SCIENCES SORBONNE UNIVERSITÉ

Balanced accuracy (↑) over 10 runs with best performance in **bold**.
Results are aggregated over imbalance scenarios and datasets.

Setting	Source	DANN	$WD_{\beta=0}$	$WD_{\beta=1}$	$WD_{\beta=2}$	MARSg	MARSc	IW-WD	OSTAR+IM
				Digits					
balanced	74.98 ± 3.8	90.81 ± 1.3	92.63 ± 1.0	82.80 ± 4.7	76.07 ± 7.1	92.18 ± 2.2	94.91 ± 1.4	95.89 ± 0.5	$\mathbf{97.51 \pm 0.3}$
subsampled	75.05 ± 3.1	89.91 ± 1.5	89.45 ± 1.0	81.56 ± 4.8	77.77 ± 6.5	91.87 ± 2.0	93.75 ± 1.4	93.22 ± 1.1	$\mathbf{96.69 \pm 0.7}$
				VisDA12					
original	48.63 ± 1.0	53.72 ± 0.9	57.40 ± 1.1	47.56 ± 0.8	36.21 ± 1.8	55.62 ± 1.6	55.33 ± 0.8	51.88 ± 1.6	$\mathbf{59.24 \pm 0.5}$
subsampled	42.46 ± 1.4	47.57 ± 0.9	47.32 ± 1.4	41.48 ± 1.6	31.83 ± 3.0	55.00 ± 1.9	51.86 ± 2.0	50.65 ± 1.5	$\mathbf{58.84 \pm 1.0}$
				Office31					
subsampled	74.50 ± 0.5	76.13 ± 0.3	76.24 ± 0.3	74.23 ± 0.5	72.40 ± 1.8	80.20 ± 0.4	80.00 ± 0.5	77.28 ± 0.4	$\mathbf{82.61 \pm 0.4}$
				OfficeHome					
subsampled	50.56 ± 2.8	50.87 ± 1.05	53.47 ± 0.7	52.24 ± 1.1	49.48 ± 1.3	56.60 ± 0.4	56.22 ± 0.6	54.87 ± 0.4	$\mathbf{59.51 \pm 0.4}$

Performance on visual UDA datasets

SCIENCES SORBONNE UNIVERSITÉ

Balanced accuracy (↑) over 10 runs with best performance in **bold**.
Results are aggregated over imbalance scenarios and datasets.

Setting	Source	DANN	$WD_{\beta=0}$	$WD_{\beta=1}$	$WD_{\beta=2}$	MARSg	MARSc	IW-WD	OSTAR+IM
					Digits				
balanced	74.98 ± 3.8	90.81 ± 1.3	92.63 ± 1.0	82.80 ± 4.7	76.07 ± 7.1	92.18 ± 2.2	94.91 ± 1.4	95.89 ± 0.5	$\mathbf{97.51 \pm 0.3}$
subsampled	75.05 ± 3.1	89.91 ± 1.5	89.45 ± 1.0	81.56 ± 4.8	77.77 ± 6.5	91.87 ± 2.0	93.75 ± 1.4	93.22 ± 1.1	$\mathbf{96.69 \pm 0.7}$
					VisDA12				
original	48.63 ± 1.0	53.72 ± 0.9	57.40 ± 1.1	47.56 ± 0.8	36.21 ± 1.8	55.62 ± 1.6	55.33 ± 0.8	51.88 ± 1.6	$\mathbf{59.24 \pm 0.5}$
subsampled	42.46 ± 1.4	47.57 ± 0.9	47.32 ± 1.4	41.48 ± 1.6	31.83 ± 3.0	55.00 ± 1.9	51.86 ± 2.0	50.65 ± 1.5	$\mathbf{58.84 \pm 1.0}$
					Office31				
subsampled	74.50 ± 0.5	76.13 ± 0.3	76.24 ± 0.3	74.23 ± 0.5	72.40 ± 1.8	80.20 ± 0.4	80.00 ± 0.5	77.28 ± 0.4	$\mathbf{82.61 \pm 0.4}$
					OfficeHome				
subsampled	50.56 ± 2.8	50.87 ± 1.05	53.47 ± 0.7	52.24 ± 1.1	49.48 ± 1.3	56.60 ± 0.4	56.22 ± 0.6	54.87 ± 0.4	$\mathbf{59.51 \pm 0.4}$

Same trends for target label estimation error $\|\boldsymbol{p}_N^Y - \boldsymbol{p}_T^Y\|_1$.

Performance on visual UDA datasets

SCIENCES SORBONNE UNIVERSITÉ

Balanced accuracy (↑) over 10 runs with best performance in **bold**.
Results are aggregated over imbalance scenarios and datasets.

Setting	Source	DANN	$WD_{\beta=0}$	$WD_{\beta=1}$	$WD_{\beta=2}$	MARSg	MARSc	IW-WD	OSTAR+IM
					Digits				
balanced	74.98 ± 3.8	90.81 ± 1.3	92.63 ± 1.0	82.80 ± 4.7	76.07 ± 7.1	92.18 ± 2.2	94.91 ± 1.4	95.89 ± 0.5	**97.51 ± 0.3**
subsampled	75.05 ± 3.1	89.91 ± 1.5	89.45 ± 1.0	81.56 ± 4.8	77.77 ± 6.5	91.87 ± 2.0	93.75 ± 1.4	93.22 ± 1.1	**96.69 ± 0.7**
					VisDA12				
original	48.63 ± 1.0	53.72 ± 0.9	57.40 ± 1.1	47.56 ± 0.8	36.21 ± 1.8	55.62 ± 1.6	55.33 ± 0.8	51.88 ± 1.6	**59.24 ± 0.5**
subsampled	42.46 ± 1.4	47.57 ± 0.9	47.32 ± 1.4	41.48 ± 1.6	31.83 ± 3.0	55.00 ± 1.9	51.86 ± 2.0	50.65 ± 1.5	**58.84 ± 1.0**
					Office31				
subsampled	74.50 ± 0.5	76.13 ± 0.3	76.24 ± 0.3	74.23 ± 0.5	72.40 ± 1.8	80.20 ± 0.4	80.00 ± 0.5	77.28 ± 0.4	**82.61 ± 0.4**
					OfficeHome				
subsampled	50.56 ± 2.8	50.87 ± 1.05	53.47 ± 0.7	52.24 ± 1.1	49.48 ± 1.3	56.60 ± 0.4	56.22 ± 0.6	54.87 ± 0.4	**59.51 ± 0.4**

Same trends for target label estimation error $\|\boldsymbol{p}_N^Y - \boldsymbol{p}_T^Y\|_1$.

Domain-invariant baselines designed for:

- Covariate shift w/o reweighting (DANN Ganin et al. 2016, $WD_{\beta=0}$ Shen et al. 2018).
- GeTarS with reweighting ($WD_{\beta\in\{1,2\}}$ Wu et al. 2019, MARS Rakotomamonjy et al. 2021; IW-WD Combes et al. 2020).

Conclusion

Conclusion

Summary

- `OSTAR`, a new general OT approach to align pretrained representations under Generalized Target Shift.
- Strong generalization guarantees under mild assumptions.

Conclusion

SCIENCES SORBONNE UNIVERSITÉ

Summary

- OSTAR, a new general OT approach to align pretrained representations under Generalized Target Shift.
- Strong generalization guarantees under mild assumptions.

Paper, code and contact

- https://openreview.net/forum?id=sPfB2PI87BZ
- https://github.com/mkirchmeyer/ostar
- Contact: matthieu.kirchmeyer@isir.upmc.fr

Conclusion

Summary

- OSTAR, a new general OT approach to align pretrained representations under Generalized Target Shift.
- Strong generalization guarantees under mild assumptions.

Paper, code and contact

- https://openreview.net/forum?id=sPfB2PI87BZ
- https://github.com/mkirchmeyer/ostar
- Contact: matthieu.kirchmeyer@isir.upmc.fr

Thank you for your attention !

References I

- Combes, Remi Tachet des et al. (2020). "Domain Adaptation with Conditional Distribution Matching and Generalized Label Shift". In: *Advances in Neural Information Processing Systems*.
- Ganin, Yaroslav et al. (2016). "Domain-Adversarial Training of Neural Networks". In: *Journal of Machine Learning Research* 17.59, pp. 1–35.
- Gong, Mingming et al. (2016). "Domain Adaptation with Conditional Transferable Components". In: *Proceedings of The 33rd ICML*. Vol. 48. Proceedings of Machine Learning Research. New York, New York, USA: PMLR, pp. 2839–2848.
- Liu, Hong et al. (2019). "Transferable Adversarial Training: A General Approach to Adapting Deep Classifiers". In: *Proceedings of 36th ICML*. Vol. 97. Proceedings of ML Research. PMLR, pp. 4013–4022.
- Rakotomamonjy, A. et al. (2021). "Optimal transport for conditional domain matching and label shift". In: *Machine Learning*.

References II

- Shen, Jian et al. (2018). "Wasserstein Distance Guided Representation Learning for Domain Adaptation". In: *AAAI-18, 30th IAAI-18, 8th AAAI Symposium on EAAI-18, New Orleans, USA, February 2-7*. AAAI Press, pp. 4058–4065.
- Shui, Changjian et al. (2021). "Aggregating From Multiple Target-Shifted Sources". In: *Proceedings of the 38th International Conference on Machine Learning*. Ed. by Marina Meila and Tong Zhang. Vol. 139. Proceedings of Machine Learning Research. PMLR, pp. 9638–9648. url: `http://proceedings.mlr.press/v139/shui21a.html`.
- Wu, Yifan et al. (2019). "Domain Adaptation with Asymmetrically-Relaxed Distribution Alignment". In: *ICML*. Vol. 97. Proceedings of ML Research. Long Beach, CA, USA: PMLR, pp. 6872–6881.
- Zhang, Kun et al. (2013). "Domain Adaptation under Target and Conditional Shift". In: *Proceedings of the 30th ICML*. Vol. 28. Proceedings of Machine Learning Research 3. Atlanta, Georgia, USA, pp. 819–827.