Visual DA in Deep Learning Era

Gabriela Csurka

NAVER LABS Europe (NLE), Meylan, France

gabriela.csurka@naverlabs.com

NAVER LABS

Outline

1. Motivation

2. Domain adaptation in Deep Learning Era

3. Deep Domain Adaptation Methods

4. Beyond image classification

Domain adaptation (DA)

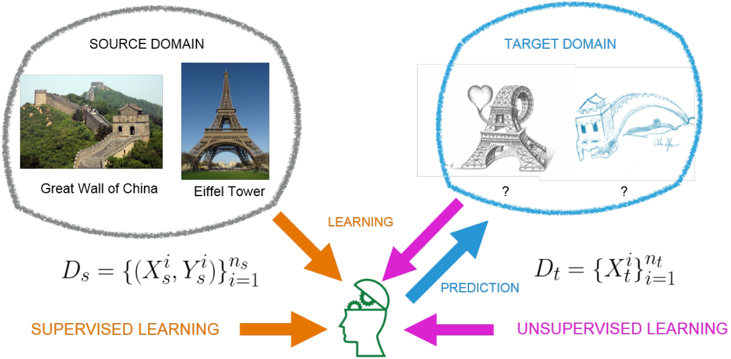

Leveraging labeled **source domain**, to learn a model for the **target domain**.

Example scenarios

Recognition

Detection

Segmentation

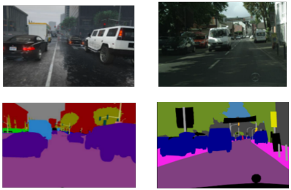

Re-identification

Control

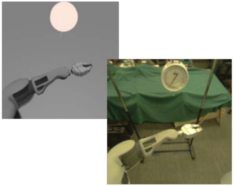

Visual localization

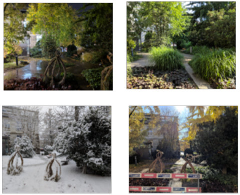

Outline

1. Motivation
2. Domain adaptation in Deep Learning Era
3. Deep Domain Adaptation Methods
4. Beyond image classification

How to exploit deep models?

Shallow methods using deep features

- use the deep model as feature extractor
- apply any shallow DA method using these features

Using fine-tuned deep architectures

- fine-tune the deep model on the source
- apply the fine-tuned model on the target

Shallow methods using fine-tuned deep features

- fine-tune the deep model on the source
- use the fine-tuned model as feature extractor
- apply any shallow DA method using these features

Deep DA models

- specific deep architectures tailored for domain adaptation
- often initialized with a deep model fine-tuned on the source

Classical Shallow DA Methods

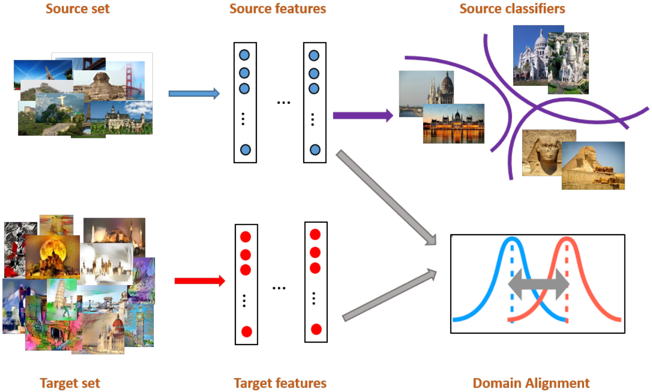

- Any pre-computed (vectorial) image representation
- Classifier: *e.g.* SVM, KNN or MLP
- Domain alignment: *e.g.* by minimizing the distribution mismatch

Shallow methods with deep features

Deep features are more abstract, already decreases the domain bias.

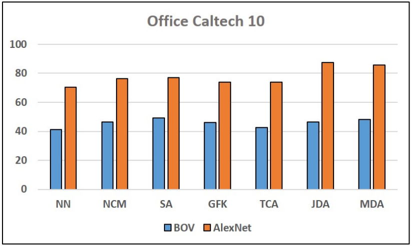

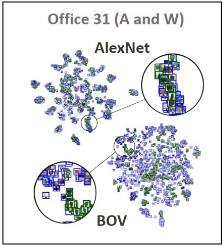

Pre-trained image classification models

- Activations layers of the deep CNN model, Donuahe+@ICML'14.

Deep image representations learning

- Trained with ranking or contrastive losses

Fine-tuning the model on the source

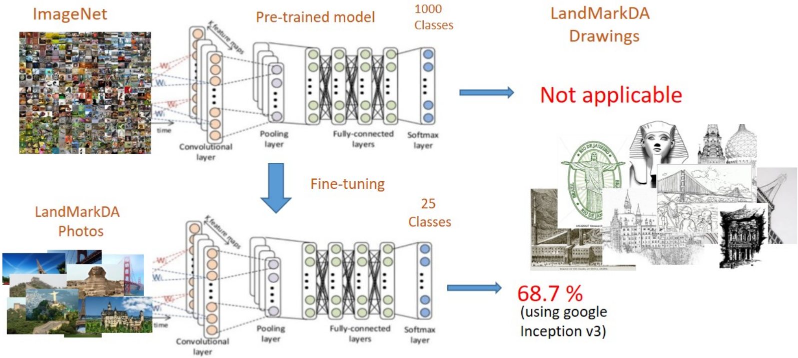

Fine-tuning the model on the source

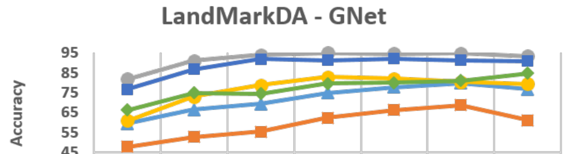

	FT(Lg)	FT(7c)	FT(7a)	FT(6e)	FT(6c)	FT(6a)	FT(all)
Ph->Pt	59.5	66.4	69.3	74.9	77.8	79.7	76.7
Ph->Dr	47.8	52.7	55.3	62.4	66.3	68.7	61.1
Pt->Ph	81.9	91	93.9	94.9	94.6	94.7	93.3
Pt->Dr	60.8	72.9	78.7	82.8	82.1	80.3	79.3
Dr->Ph	76.6	86.8	91.9	91.2	92	91.3	90.8
Dr->Pt	66.3	74.8	74.4	79.5	79.9	80.9	84.7

- Fine-tuning deeper is better than finetuning only the last layers
- How deep we need to fine-tune the model depends on the domain gap

Deep versus Shallow models

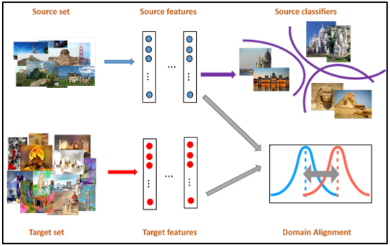

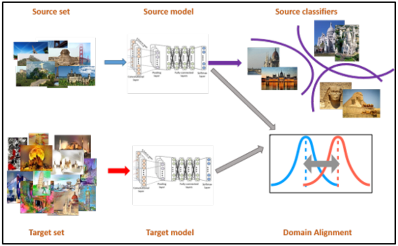

Shallow models:

- acts on pre-extracted (deep) image representations
- learns independently or jointly the latent space and the classifiers

Deep models:

- acts directly on the images
- learns image representation, domain bias and the classifier all end-to-end

Deep DA model versus deep features

DeepCORAL vs CORAL, Sun+@TASK-CV'16

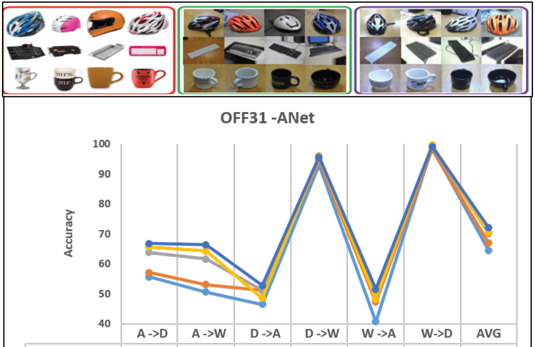

	A ->D	A ->W	D ->A	D ->W	W ->A	W->D	AVG
NA	55.7	50.6	46.5	93.1	40.7	98.9	64.4
CORAL	57.1	53.1	51.1	94.6	47.3	98.2	66.9
FT	63.8	61.6	51.1	95.4	49.8	99	70.1
FT+ CORAL	65.7	64.3	48.5	96.1	48.2	99.8	70.4
D-CORAL	66.8	66.4	52.8	95.7	51.5	99.2	72.1

- Shallow model improves little over directly using deep feature.

Deep DA model versus deep features

Discrepancy-based deep vs shallow networks, Csurka[+]@TASK-CV'17

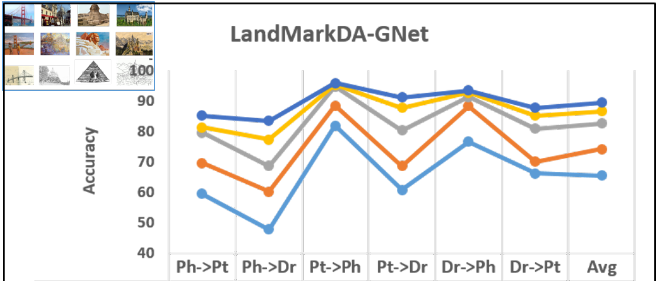

	Ph->Pt	Ph->Dr	Pt->Ph	Pt->Dr	Dr->Ph	Dr->Pt	Avg
NA	59.5	47.8	81.9	60.8	76.6	66.3	65.5
SDAN	69.6	60.3	88.5	68.6	88.3	70.1	74.2
FT	79.7	68.7	94.7	80.3	91.3	80.9	82.6
FT + SDAN	81.3	77.4	95.3	87.8	92.7	85.2	86.6
DDAN	85.2	83.5	95.9	91.2	93.4	87.7	89.5

- Fine-tuning the deep model on the source outperforms the shallow model.
- Shallow with fine-tuned deep features is close to deep model (the best).

To summarize

Shallow models with deep features

- simple and low cost solutions
- same architecture can be applied to any vectorial representation

Tailored deep DA models

- can adjust the feature representation to the problem
- if appropriately trained they often outperform the shallow methods

Shallow methods using fine-tuned deep features

- combines the strength of deep learning and domain adaptation
- fine-tuning can be done in advance, before seeing the target
- no need for new complex architecture
- close to results obtained with the adapted DA model

Outline

1. Motivation
2. Domain adaptation in Deep Learning Era
3. Deep Domain Adaptation Methods
4. Beyond image classification

Discriminative models

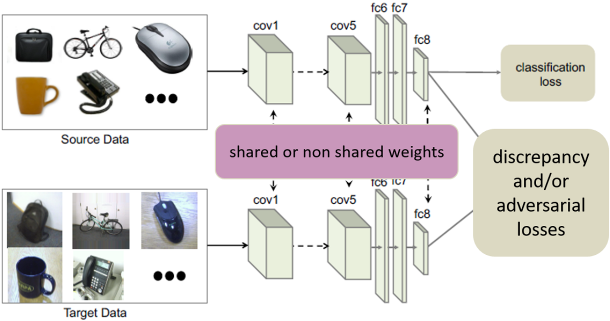

- Siamese network, one source and one target stream
  - Both stream initialised with the pretrained-model on the source
- Classification (cross-entropy) loss on the source
- Domain alignment:
  - minimizing the distribution discrepancy
  - adversarial domain confusion

Minimizing feature distribution discrepancy

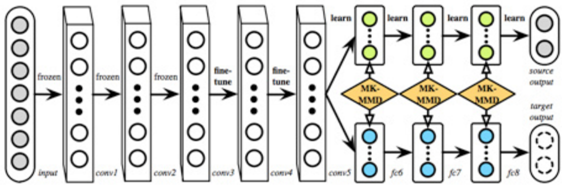

- Kernelized MMD loss, DAN (Long$^+$@ICML'15)

$$MMD(S,T) = \sum_{l=1}^{L} \| \mathbb{E}(\phi(M_S^l)) - \mathbb{E}(\phi(M_T^l)) \|_2$$

where ϕ is a kernel projection and $\mathbb{E}(X) = \frac{1}{|X|} \sum_{x \in X}$ is the empirical expectation.

- Weighted discrepency, WDAN (Yan$^+$@CVPR'17)

Alternative discrepancy losses

- Central Moment Discrepency, CMD (Zellinger$^+$@ICLR'17)

$$CMD(S,T) = \| \mathbb{E}(M_S) - \mathbb{E}(M_T) \|_2 + \sum_{k=2}^{\infty} \frac{1}{|b-a|^k} \| C_k(M_S) - C_k(M_T) \|$$

where $C_k(X) = \mathbb{E}((x - \mathbb{E}(X))^k)$ is the k^{th} order sample central moment.

- Wasserstein Distance: WGRL (Shen$^+$@AAAI'18), NWD (Balaji$^+$@ICCV'19)

$$WD(S,T) = \sup_{\|\phi\|_{L\leq 1}} \left(\mathbb{E}_{P_S} \left[\phi(M_S(\mathbf{x}^S)) \right] - \mathbb{E}_{P_T} \left[\phi(M_T(\mathbf{x}^T)) \right] \right)$$

where $\| \cdot \|_L$ is the Lipschitz semi-norm, P_S and P_T are marginal distributions.

- Deep correlation alignement, DeepCORAL (Sun$^+$@TASK-CV'16)

$$CORAL(S,T) = \frac{1}{4d^2} \| Cov(M_S)) - Cov(M_T)) \|_F^2$$

where $Cov(X)$ is the data covariance of X.

Adversarial learning

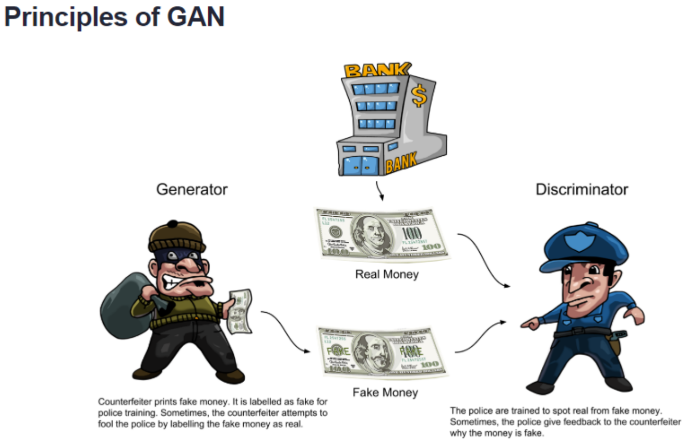

Image: Courtesy to Richard Gall.

- Generative adversarial nets (GAN), Goodfellow$^+$@NIPS'14

Domain adversarial training

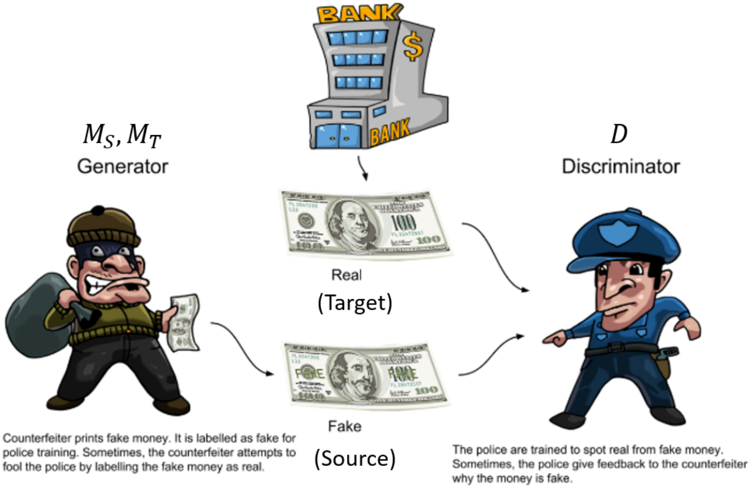

Increase domain confusion

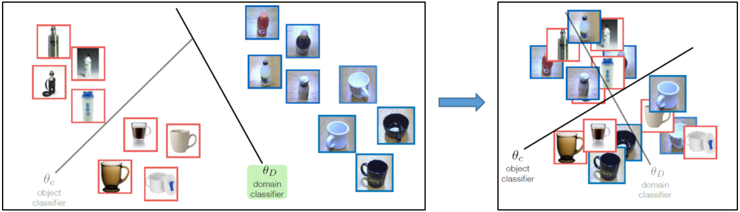

Image: Courtesy to Judy Hoffman.

- Adversarial (GAN) loss, ADDA (Tzeng+@CVPR'17)

$$\max_{D} \{\mathbb{E}_{\mathbf{x}\sim p_S(\mathbf{x})}[\log D(M_S(\mathbf{x}))] + \mathbb{E}_{\mathbf{x}\sim p_T(\mathbf{x})}[\log(1 - D(M_T(\mathbf{x}))]\}$$

$$\max_{M_T} \{\mathbb{E}_{\mathbf{x}\sim p_T(\mathbf{x})}[\log D(M_T(\mathbf{x}))]\}$$

- Deep domain confusion, DDC (Tzeng+@ARXIV'14)
- Jensen-Shannon divergence (by GAN), GAM (Huang+@ECCV'18)

Gradient reversal layers

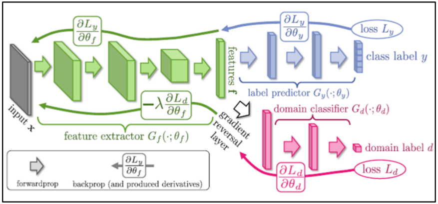

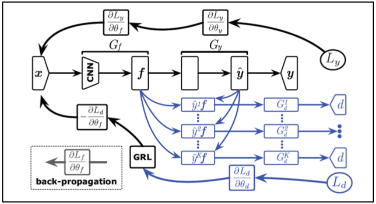

- RevGrad (Ganin+@JMLR'16), MADA (Pei+@AAAI'18), SimNet (Pinhero+@CVPR'18)

$$\min_{M_S, M_T} \max_{D} V(D, M_S, M_T) = \mathbb{E}_{\mathbf{x}\sim p_S(\mathbf{x})}[\log D(M_S(\mathbf{x}))]\mathbb{E}_{\mathbf{x}\sim p_T(\mathbf{x})}[\log(1 - D(M_T(\mathbf{x})))]$$

Experimental comparisons

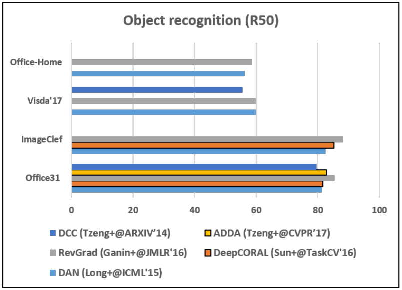

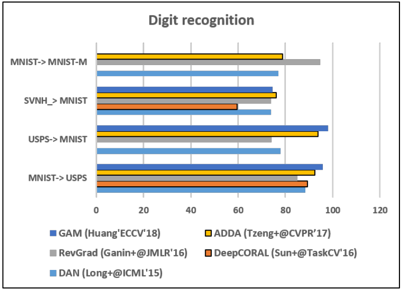

- Adversarial losses (ADDA, RevGrad, GAM) performs in general better than feature discrepency minimization (DAN, DeepCORAL).

Target network parameter adaptation

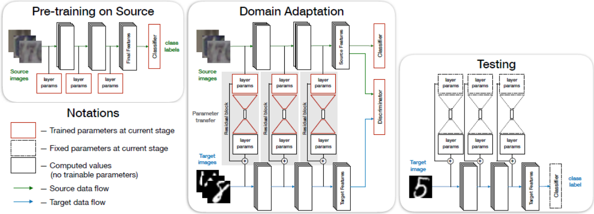

- Linear global transformations, BSW (Rozantsev+@PAMI'18)

$$r_w(\theta^l_S, \theta^l_T) = \exp\left(\|a_l\theta^l_S + b_l - \theta^l_T\|^2\right) - 1$$

where a_l and b_l are scalars learned during the training.

- Residual parameter transfer, RPT (Rozantsev+@CVPR'18)

$$\Theta^l_t - \theta^l_S = \mathbf{B}^l_1 \sigma\left((\mathbf{A}^l_1 \theta^l_S \mathbf{A}^l_2 + \mathbf{D}^l)\mathbf{B}^l_2\right)$$

where $\mathbf{A}^l_1, \mathbf{A}^l_2, \mathbf{B}^l_1, \mathbf{D}^l_2, \mathbf{D}^l$ are transformation parameters at layer l.

Target network parameter adaptation

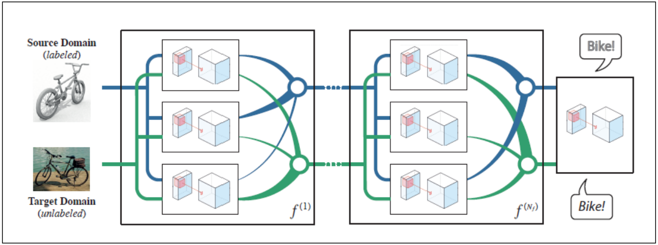

- Domain adaptive multi-branch network, DAMNet (Bermúdez-Chacón+@ICLR'19)

$$x^l = \sum_k a_k^l \theta_k^l (x^{l-1})$$

where a_k^l are the trainable activation weights of the gates.

Experimental comparisons

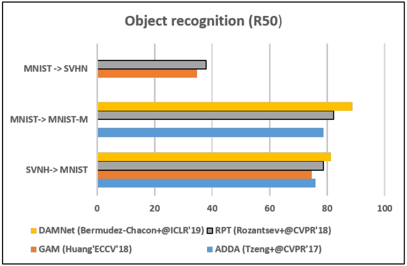

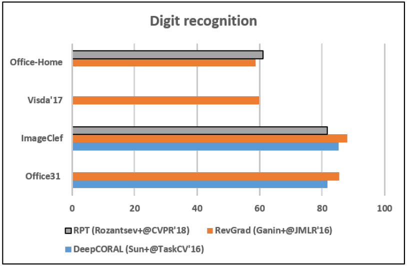

- Best strategy seems to be the gated multi-branch network (DAMNet)

Adapting the batch

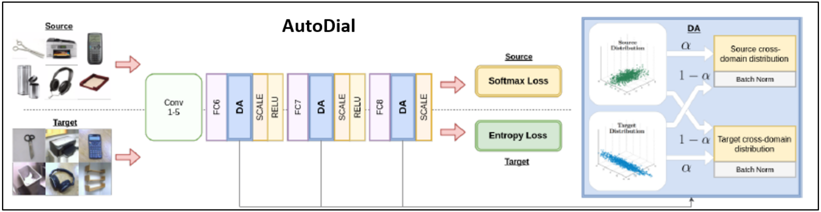

- Domain specific batch normalization, AutoDial (Carlucci+@ICCV'17), AdaBN (Li+@PR'18), DSBN (Chang+@CVPR'19)
- Batch Nuclear-norm Maximization, BNM (Cui+@CVPR'20)
- Batch Whitening, DWT (Roy+@CVPR'19)
- Learning batch re-weighting with mass shift, JD-BW (Binkowski+@ICCV'19)

Experimental comparisons

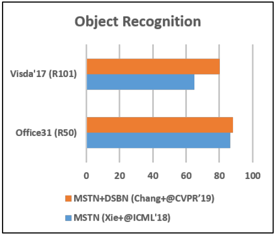

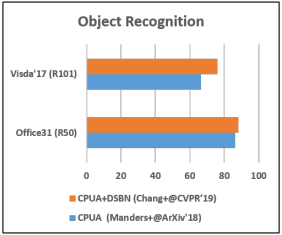

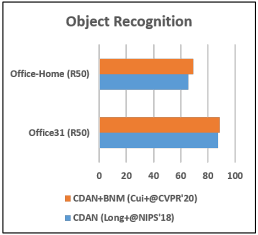

- Adapting batch normalization for the target helps (DSBN, BNM).

Transfer domain style

Paired I2I

Un-paired I2I

Paired image-to-image style transfer as preprocessing

- Csurka$^+$@TASKCV'17, Thomas$^+$@ACCV'19, Jackson$^+$@CVPR-WS'19

Unpaired image-to-image style transfer learning

- I2I (Zhu$^+$@ICCV'17), I2IAd (Murez$^+$@CVPR'18)

Transfer domain style with GAN

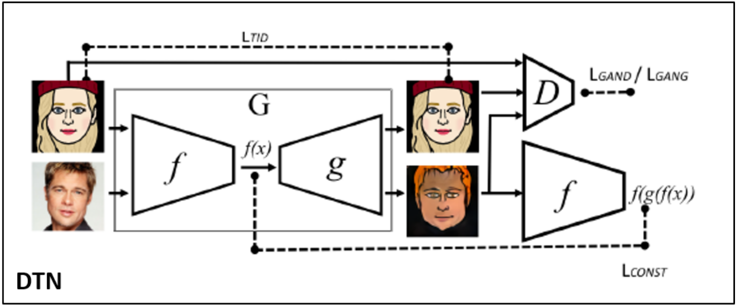

Single GAN

- PLDT (Yoo+@ECCV'16), PixelDA (Bousmalis+@CVPR'17), DTN (Taigman+@ICLR'17), GenToAdapt (Sankaranarayanan+@CVPR'18)

Combine several GANs

- CoGAN (Liu+@NIPS'16), UNIT (Liu+@NIPS'17), DupGAN (Hu+@CVPR'18)

Align images (CycleGAN) and image representations

- CyCADA (Hoffman+@ICML'18), DRIT (Lee+@ECCV'18), ContrAN (Kang+@CVPR'19)

Encoder-decoder based models

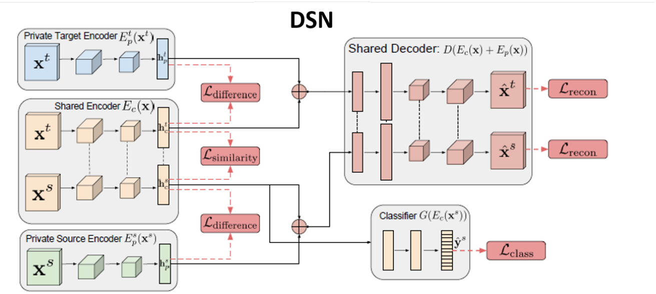

Shared encoder-decoder

- sMDA, Chen$^+$@ICML'12, TLDA, Zhuang$^+$@IJCAI'15

Domain specific encoding and/or decoding

- DRCN, Ghifary$^+$@ECCV'16, DSN, Bousmalis$^+$@NIPS'16

Experimental comparisons

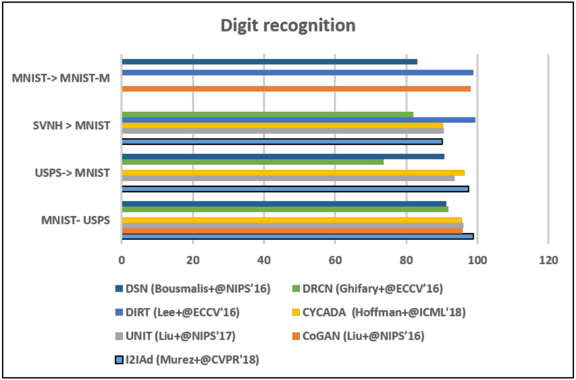

- Adversarial I2I transformation performs better than unsupervised encoder-decoder based reconstruction (DSN, DRCN)
- Best results obtained when both the images and their representation are aligned (I2IAd, CyCADA, DRIT)

Consistency between multiple source

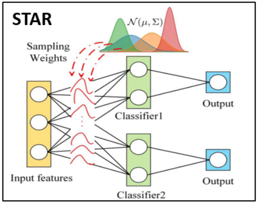

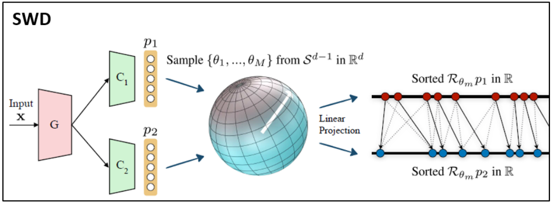

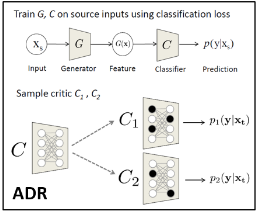

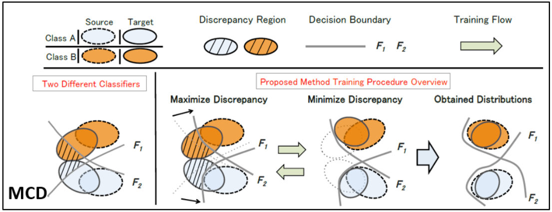

Diversify source classifier

- MCD (Saito⁺@CVPR'19), ADR (Saito⁺@ICLR'18), SWD (Lee⁺@CVPR'19), STAR (Lu⁺@CVPR'20)

Cyclic consistency

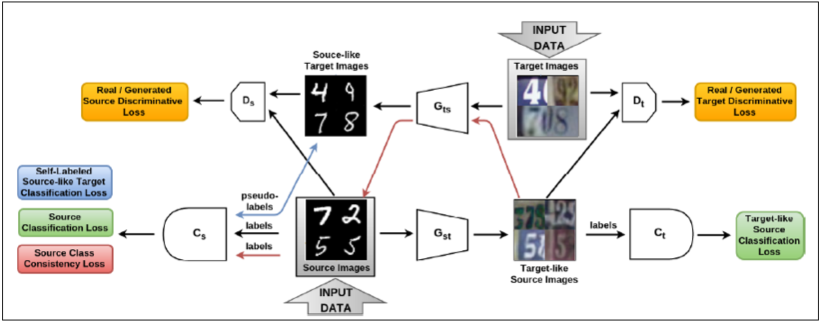

- Predict source from predicted target, LTR (Sener+@NIPS'16)
- Predict from traget-like source image, SBADA-GAN (Russo+@ARXIV'17)

Experimental comparisons

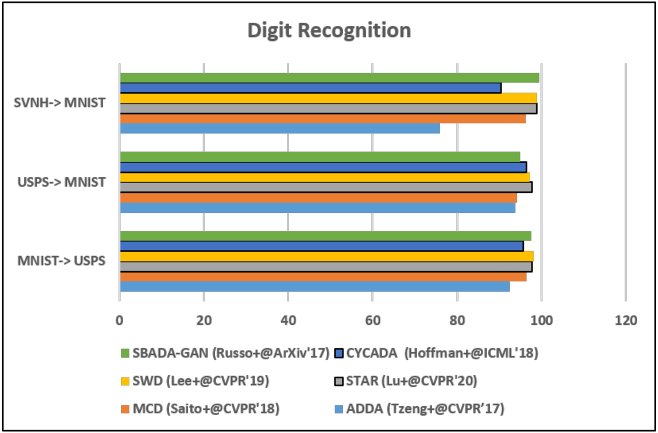

- Significant improvement over the corresponding baseline methods.

Deep optimal transport

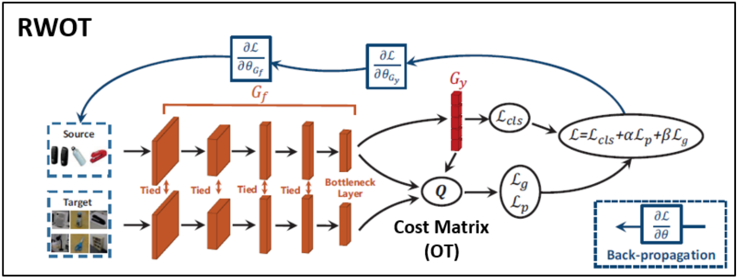

Source class information guides the optimal transport

- DeepJDOT (Damodaran+@ECCV'16), RWOT (Xu+@CVPR'20)

Joint feature and score distributions

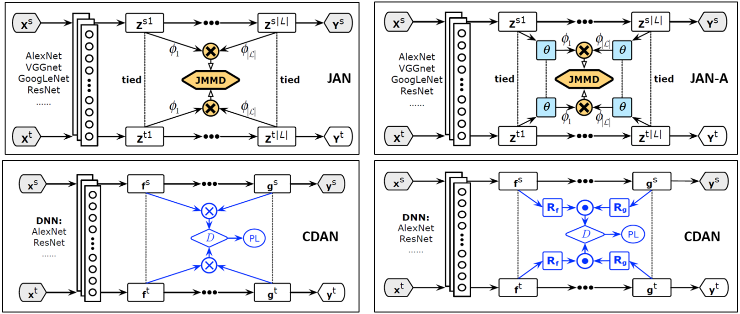

- Joint distribution alignment, JAN (Long^{+}@ICML'17)
- Adversarial joint adaptation, JAN-A (Long^{+}@NIPS'18)
- Conditional domain adversarial network, CDAN (Long^{+}@NIPS'18)

Experimental comparisons

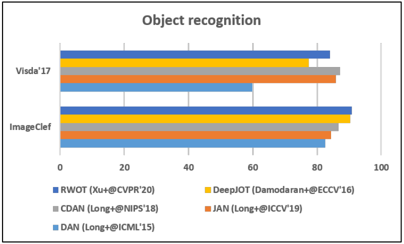

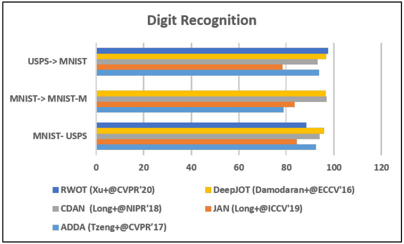

- Best overall CDAN (Long$^+$@NIPS'18) and DeepJDOT (Damodaran$^+$@ECCV'16)

Target score distribution entropy

Minimize the entropy of the target predictions (MinEnt)

- AutoDial (Carlucci+@ICCV'17), ATT (Saito+@ICML'17), SBADA-GAN (Russo+@ARXIV'17), DTA (Lee+@ICCV'19), RCA (Cicek+@ICCV'19)

$$\sum_{x^T} \sum_{y \in \mathcal{C}} p(y|x^T) \log p(y|x^T)$$

Min-Entropy Consensus (MEC)

- DWT-MEC (Roy+@CVPR'19)

$$-\frac{1}{2} \sum_{x^T} \max_{y \in \mathcal{C}} \left(\log p(y|x_1^T) + \log p(y|x_2^T) \right)$$

where x_1^T and x_2^T are two perturbed versions of x^T.

Adversarial, Min-Max Entropy (MME)

- MME (Saito+@CVPR'19)

$$\theta_F^* = \mathrm{argmin}_{\theta_F} + \lambda H \quad \text{and} \quad \theta_C^* = \mathrm{argmin}_{\theta_C} - \lambda H$$

where H is the entropy, and θ_F, θ_C are the parameters of the feature extractor and classifier respectively.

Experimental comparisons

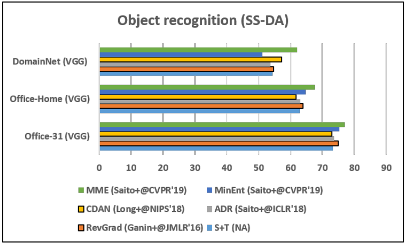

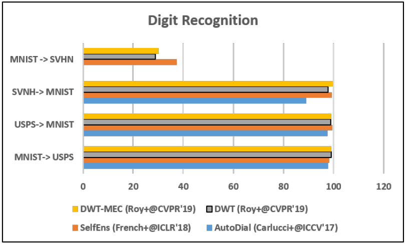

- MEC and MME seems to be better than using simply MinEnt.

Teacher-student paradigm

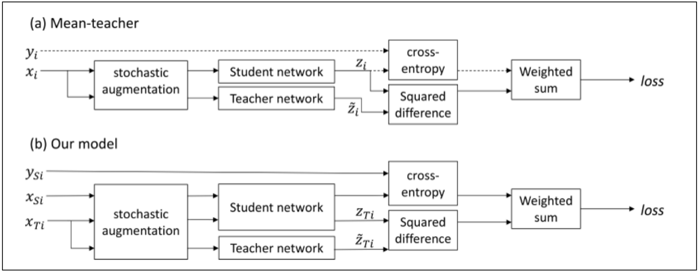

Mean-teacher of data augmented ensemble classifier

- SelfEns (French+@ICLR'18), DWT (Roy+@CVPR'19)

Refine student classifier's decision-boundary with a teacher

- DIRT-T (Shu+@ICLR'18)

Cluster alignment with a teacher

- CAT (Deng, +@ICCV'19)

Experimental comparisons

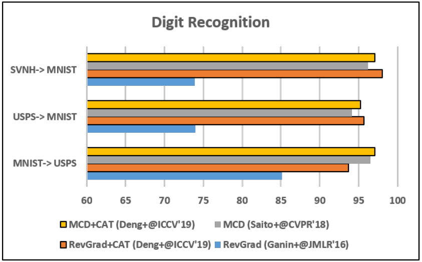

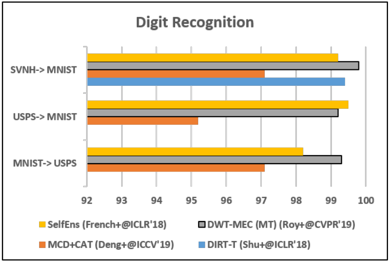

- Adding CAT improves the corresponding model baseline model.
- Mean Teacher of ensemble classifier performs the best.

Curriculum/Self-learning

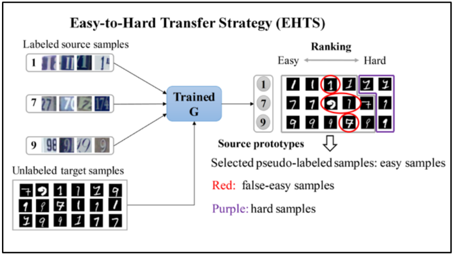

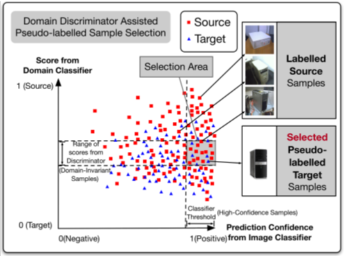

Easy-to-hard sample strategy (ETHS)

- PFAN (Chen+@CVPR'19)

Select highly confident and domain uninformative examples

- iCAN (Zhang+@CVPR'18)

Curriculum based dropout discriminator

- CD^{3}A (Kurmi, +@BMVC'19)

Contrastive intra and inter-class domain discrepancy optimization

- ContrAN (Kang+@CVPR'19)

Experimental comparisons

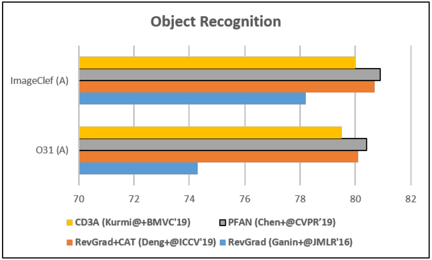

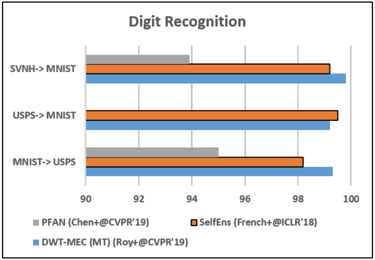

- Easy-to-hard sample strategy (PFAN) seems to be the best on object recognition.
- Performs less well as the ensemble learning (SelfEns, DWT) on digit recognition.

To summarize

Winning strategies:

- Adversarial adaptation vs discriminative (CDAN, GAM)
- GAN (CyCADA, DRIT) better vs encoder-decoder
- Exploit score distributions to guide feature alignment (MCD, RWOT, DWT)
- Curriculum/Self-learning using pseudo-labels (PFAN, iCAN)

The results are to be taken cautiously as

- The results come from various papers
- Not clear how the hyperparameters for each model were selected
- Not always clear how comparable the models (*e.g*. diff architecture)

Best result on object recognition sets

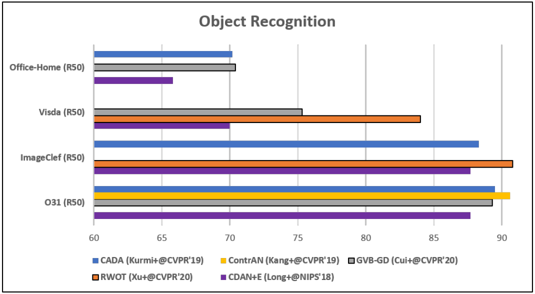

Best results are often complex models which in general

- acts mainly on the score prediction level
- exploit target prediction uncertainty
- and has some specific ingredient

Best result on object recognition sets

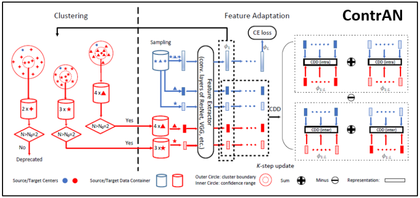

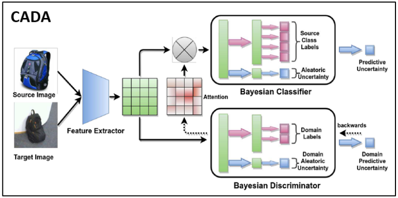

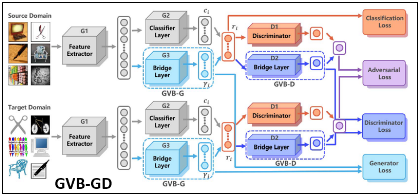

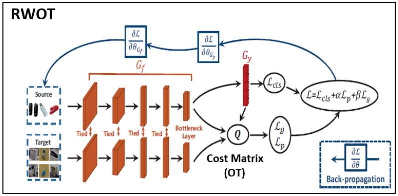

- Contrastive discrepancy adaptation, ContrAN (Kang^{+}@CVPR’19)
- Uncertainty based attention, CADA (Kurmi^{+}@CVPR’19)
- Gradually vanishing bridges, GVB-GD (Cui^{+}@CVPR’20)
- Weighted optimal transport, RWOT (Xu^{+}@CVPR’20)

Outline

1. Motivation
2. Domain adaptation in Deep Learning Era
3. Deep Domain Adaptation Methods
4. Beyond image classification

DeepDA becoming extremely popular in CV

Many method proposed for:

- Semantic segmentation
- Person Re-ID
- Object detection

But recent DA methods were also proposed for:

- Pose/action recognition
- Depth estimation
- Low level image enhancement
- Control in robotics
- 3D/Visual localization
- Medical imaging

Image Segmentation

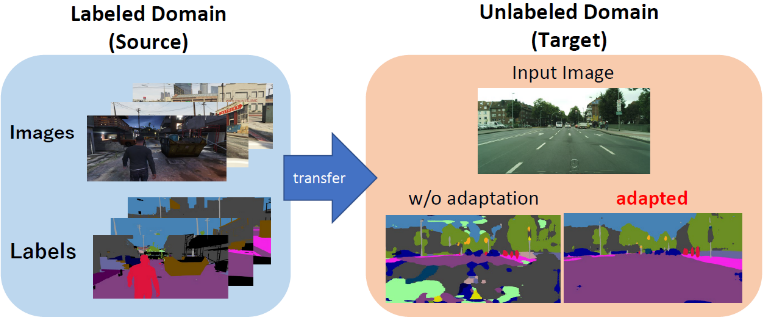

From Synthetic to real data

- Easy to obtain pixel level annotation
- Poor labeling due to domain shift

Segmentation model adaptation

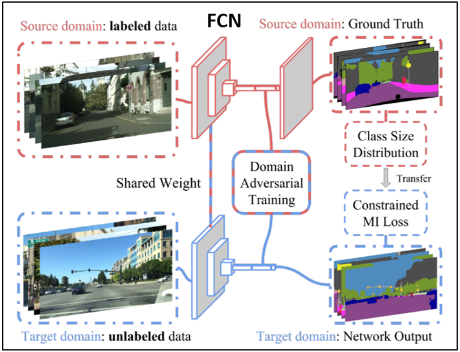

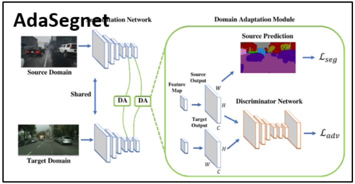

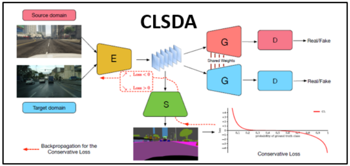

- Transferring label statistics, FCN-WLD (Hoffman+@CORR'16)
- Backpropagating contrastive loss, CLSDA (Zhu+@ECCV'18)
- Multilevel Adversarial Learning, AdaSegNet (Tsai+@CVPR'18)

Appearance adaptation

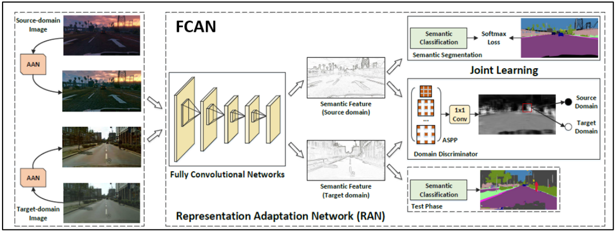

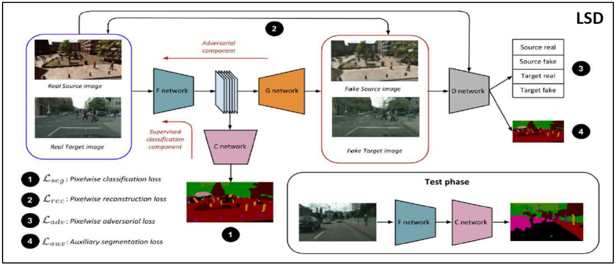

- Paired style transfer, FCAN (Zhang$^+$@CVPR'18)
- GAN based (unpaired), GenToAdapt (Sankaranarayanan$^+$@CVPR'18)

Cyclic I2I transfer

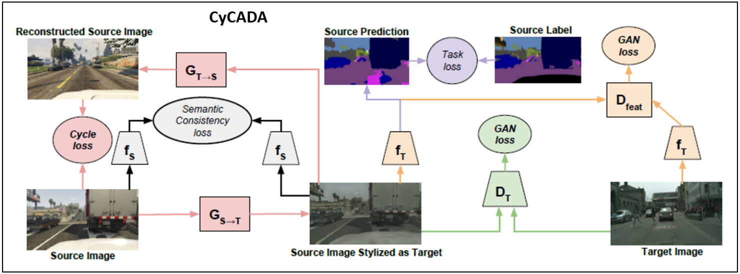

- Segmentation consistency, CyCADA (Hoffman+@ICML'18)
- Domain agnostic latent space, I2IT (Murez+@CVPR'18)
- Dual channel-wise feature alignment, DCAN (Wu+@ECCV'18)
- Cross-domain consistency, CroDoCo (Chen+@CVPR'19)
- Domain-invariant structure extraction, DISE (Chang+@CVPR'19)

Multiple source classifier

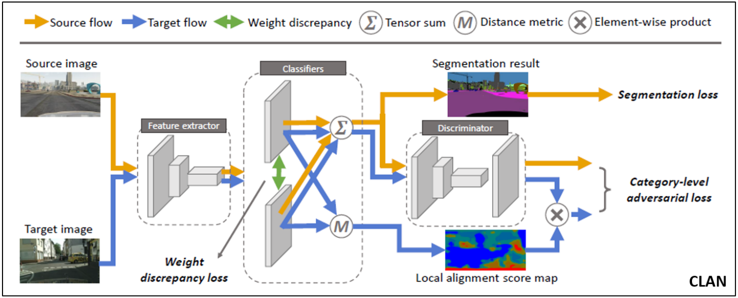

- Sliced Wasserstein Discrepancy, SWD (Lee+@CVPR'17)
- Classifiers consensus maximization, MCD (Saito+@CVPR'18)
- Minimizing the cosine similarity, CLAN (Luo+@CVPR'19)

Self-training learning

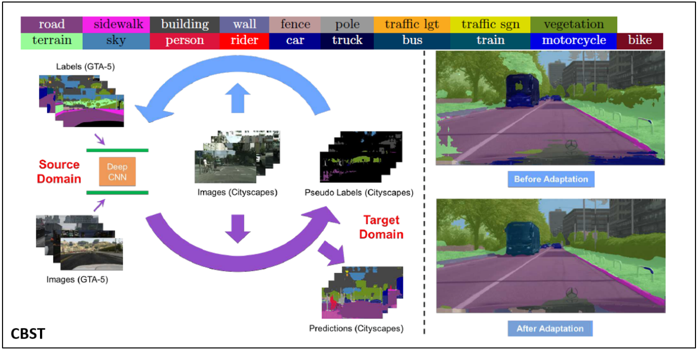

- Class-balanced self-training, CBST (Zou^{+}@ECCV'18)
- Bidirectional learning, BDL (Li^{+}@CVPR'19)
- Differential Treatment for Stuff and Things, SIM (Wang^{+}@CVPR'20)

Exploiting the prediction entropy/confidence

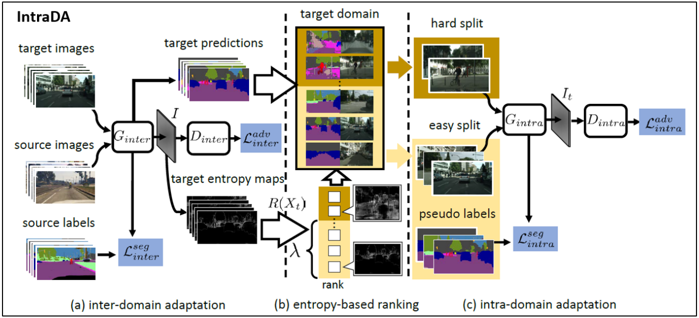

- Adversarial entropy minimization, AdvEnt (Vu+@CVPR'19)
- Progressive confidence based reweighting, SSF-DAN (Du+@ICCV'19)
- Maximum Squares Loss, MSL (Chen+@ICCV'19)
- Fourier Domain Adaptation, FDA (Yang+@CVPR'20)
- Intra-domain Adaptation, IntraDA, (Pan+@CVPR'20)

Curriculum learning

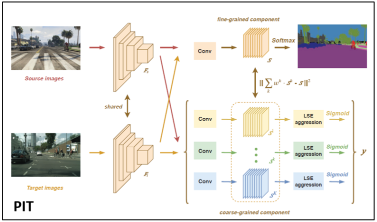

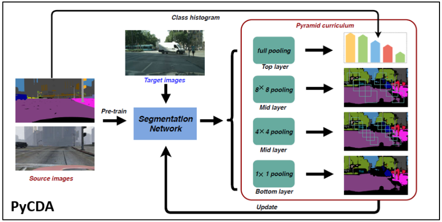

- Using static object prior, CrossCity (Chen$^+$@ICCV'17)
- Inferring first label distributions for image and landmark superpixels, CDA (Zhang$^+$@ICCV'17)
- Pyramid curriculum domain adaptation, PyCDA (Lian$^+$@ICCV'19)
- Course-to-fine region expansion, PIT (Lv$^+$@CVPR'20)

Learning with Privileged Information

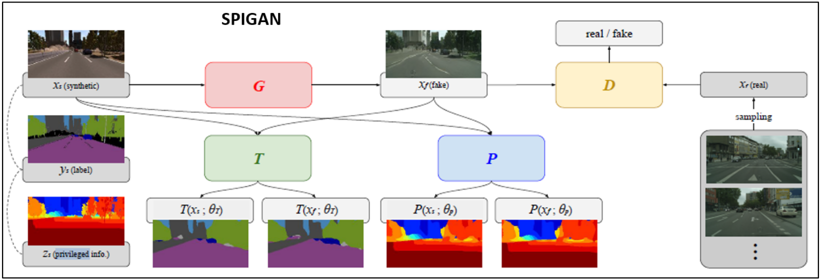

- Depth as auxiliary task, SPIGAN (Lee$^+$@ICLR'19), DADA (Vu$^+$@ICCV'19)

Cross-Modal 2D-3D segmentation

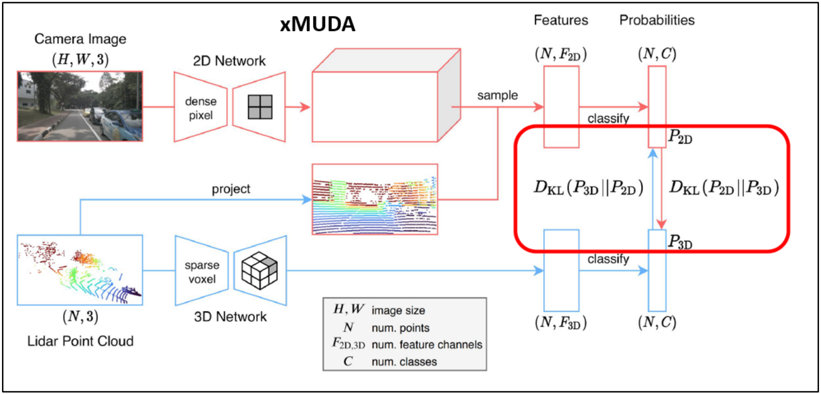

- Joint 2D image and 3D point clouds segmentation xMUDA (Jaritz+@CVPR'20)

Object detection

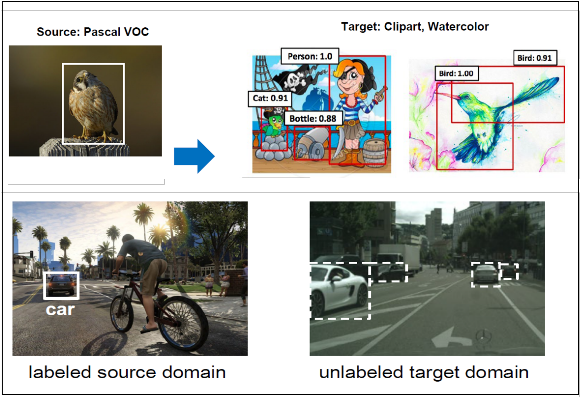

- Adapting Faster R-CNN, Chen^{+}@CVPR'18, Zhu^{+}@CVPR'19, Saito^{+}@CVPR'19, Xu^{+}@CVPR'20
- Self-training, RoyChowdhury^{+}@CVPR'19, Inoue^{+}@CVPR'18, Kim^{+}@ICCV'19

Other Visual Applications

Person Re-ID

- Wei+@CVPR'18, Liu+@CVPR'18, Zhong+@CVPR'18, Bak+@ECCV'18, Song+@CVPR'19, Fu+@ICCV'19, QI+@ICCV'19, Zhai+@CVPR'20, Luo+@CVPR'20

Pose/action recognition

- Yusuf+@BMVC'18, Perrett+@CVPR'19, Cao+@ICCV'19, Kuhnke+@ICCV'19, Munro+@CVPR'20

Depth estimation

- Kundu+@CVPR'18, Atapour-Abarghouei+@CVPR'18, Zhao+@CVPR'20, Chidlovskii+@TASK-CV'20

Low level image analyses

- Agresti+@CVPR'19, Lu+@CVPR'19, Lin+@CVPR'19, Yan+@CVPR'20, Usman+@ICCV'19

Control in robotics

- Yang+@ECCV'18, James+@CVPR'19, Wulfmeier+@IROS'17, Tobin+@IROS'17

3D/Visual localization

- Zhou+@ECCV'18, Larsson+@ICCV'19,Piao+@ICCV'19

Medical imaging

- Bermúdez-Chacón+@ISBI'18, Perone+@NEUROIMAGE'19, Dong+@CVPR'20

References -1-

- Agresti+@CVPR'19 - Unsupervised Domain Adaptation for ToF Data Denoising With Adversarial Learning
- Atapour-Abarghouei+@CVPR'18 - Real-Time Monocular Depth Estimation Using Synthetic Data With Domain Adaptation via Image Style Transfer
- Bak+@ECCV'18 - Domain Adaptation through Synthesis for Unsupervised Person Re-identification
- Balaji+@ICCV'19 - Normalized Wasserstein Distance for Mixture Distributions with Applications in Adversarial Learning and Domain Adaptation
- Bermúdez-Chacón+@ISBI'18 - A domain-adaptive two-stream U-Net for electron microscopy image segmentation
- Bermúdez-Chacón+@ICLR'19 - Domain Adaptive Multibranch Networks
- Binkowski+@ICCV'19 - Batch weight for domain adaptation with mass shift
- Bousmalis+@NIPS'16 - Domain Separation Networks
- Bousmalis+@CVPR'17 - Unsupervised Pixel-Level Domain Adaptation with Generative Adversarial Networks
- Carlucci+@ICCV'17 - AutoDIAL: Automatic DomaIn Alignment Layers
- Cao+@ICCV'19 - Cross-Domain Adaptation for Animal Pose Estimation
- Chang+@CVPR'19 - All About Structure: Adapting Structural Information Across Domains for Boosting Semantic Segmentation
- Chang+@CVPR'19 - Domain-Specific Batch Normalization for Unsupervised Domain Adaptation
- Chen+@ICML'12 - Marginalized Denoising Autoencoders for Domain Adaptation
- Chen+@ICCV'17 - No More Discrimination: Cross City Adaptation of Road Scene Segmenters
- Chen+@CVPR'18 - Re-Weighted Adversarial Adaptation Network for Unsupervised Domain Adaptation
- Chen+@CVPR'18 - Domain Adaptive Faster R-CNN for Object Detection in the Wild
- Chen+@ICCV'19 - Domain Adaptation for Semantic Segmentation With Maximum Squares Loss
- Chen+@CVPR'19 - Progressive Feature Alignment for Unsupervised Domain Adaptation
- Chen+@CVPR'19 -CrDoCo: Pixel-Level Domain Transfer With Cross-Domain Consistency
- Chidlovskii+@TASK-CV'20 - Adversarial Transfer of Pose Estimation Regression

References -2-

- Cicek+@ICCV'19 - Unsupervised Domain Adaptation via Regularized Conditional Alignment
- Csurka+@TASKCV'17 - Discrepancy-Based Networks for Unsupervised Domain Adaptation: A Comparative Study
- Cui+@CVPR'20 - Gradually Vanishing Bridge for Adversarial Domain Adaptation
- Cui+@CVPR'20 - Towards Discriminability and Diversity: Batch Nuclear-Norm Maximization Under Label Insufficient Situations
- Damodaran+@ECCV'16 - DeepJDOT: Deep Joint Distribution Optimal Transport for Unsupervised Domain Adaptation
- Deng, +@ICCV'19 - Cluster Alignment with a Teacher for Unsupervised Domain Adaptation
- Dong+@CVPR'20 - What Can Be Transferred: Unsupervised Domain Adaptation for Endoscopic Lesions Segmentation
- Donuahe+@ICML'14 - DeCAF: A Deep Convolutional Activation Feature for Generic Visual Recognition
- Du+@ICCV'19 - SSF-DAN: Separated semantic feature based domain adaptation network for semantic segmentation
- French+@ICLR'18 - Self-ensembling for Visual Domain Adaptation
- Fuk+@ICCV'19 - Yang Fu, Yunchao Wei, Guanshuo Wang, Yuqian Zhou, Honghui Shi, Thomas S. Huang
- Ganin+@JMLR'16 - Domain-Adversarial Training of Neural Networks
- Ghifary+@ECCV'16 - Deep Reconstruction-Classification Networks for Unsupervised Domain Adaptation
- Goodfellow+@NIPS'14 - Generative adversarial nets
- Hoffman+@CORR'16 - FCNs in the Wild: Pixel-level Adversarial and Constraint-based Adaptation
- Hoffman+@ICML'18 - CyCADA: Cycle-Consistent Adversarial Domain Adaptation
- Hu+@CVPR'18 - Duplex Generative Adversarial Network for Unsupervised Domain Adaptation
- Huang+@ECCV'18 - Domain Transfer Through Deep Activation Matching
- Inoue+@CVPR'18 - Cross-domain weakly-supervised object detection through progressive domain adaptation
- Jackson+@CVPR-WS'19 - Style Augmentation: Data Augmentation via Style Randomization
- James+@CVPR'19 - Sim-To-Real via Sim-To-Sim: Data-Efficient Robotic Grasping via Randomized-To-Canonical Adaptation Networks
- Jaritz+@CVPR'20 - xMUDA: Cross-Modal Unsupervised Domain Adaptation for 3D Semantic Segmentation

References -3-

- Kang+@CVPR'19 - Contrastive Adaptation Network for Unsupervised Domain Adaptation
- Kim+@ICCV'19 - Self-training and adversarial background regularization for unsupervised domain adaptive one-stage object detection.
- Kuhnke+@ICCV'19 - Deep Head Pose Estimation Using Synthetic Images and Partial Adversarial Domain Adaption for Continuous Label Spaces
- Kundu+@CVPR'18 - AdaDepth: Unsupervised Content Congruent Adaptation for Depth Estimation
- Kurmi+@BMVC'19 - Curriculum based Dropout Discriminator for Domain Adaptation
- Kurmi+@CVPR'19 - Attending to Discriminative Certainty for Domain Adaptation
- Larsson+@ICCV'19 - Fine-Grained Segmentation Networks: Self-Supervised Segmentation for Improved Long-Term Visual Localization
- Lee+@ECCV'18 - Diverse Image-to-Image Translation via Disentangled Representations
- Lee+@ICLR'19 - SPIGAN: Privileged Adversarial Learning from Simulation
- Lee+@CVPR'19 - Sliced Wasserstein Discrepancy for Unsupervised Domain Adaptation
- Lee+@ICCV'19 - Drop to Adapt: Learning Discriminative Features for Unsupervised Domain Adaptation
- Lee+@CVPR'20 - SPIGAN: Privileged Adversarial Learning from Simulation
- Li+@PR'18 - Revisiting Batch Normalization For Practical Domain Adaptation
- Li+@CVPR'19 - Bidirectional Learning for Domain Adaptation of Semantic Segmentation
- Lian+@ICCV'19 - Constructing self-motivated pyramid curriculums for crossdomain semantic segmentation: A non-adversarial approach
- Lin+@CVPR'19 - DuDoNet: Dual Domain Network for CT Metal Artifact Reduction
- Liu+@NIPS'16 - Coupled Generative Adversarial Networks
- Liu+@NIPS'17 - Unsupervised Image-to-Image Translation Networks
- Liu+@CVPR'18 - Detach and Adapt: Learning Cross-Domain Disentangled Deep Representation
- Long+@ICML'15 - Learning Transferable Features with Deep Adaptation Networks
- Long+@ICML'17 - Deep Transfer Learning with Joint Adaptation Networks
- Long+@NIPS'18 - Conditional Adversarial Domain Adaptation

References -4-

- Lu+@CVPR'19 - Unsupervised Domain-Specific Deblurring via Disentangled Representations
- Lu+@CVPR'20 - Stochastic Classifiers for Unsupervised Domain Adaptation
- Luo+@CVPR'19 -Taking a Closer Look at Domain Shift: Category-Level Adversaries for Semantics Consistent Domain Adaptation
- Luo+@ECCV'20 - Generalizing Person Re-Identification by Camera-Aware Invariance Learning and Cross-Domain Mixup
- Lv+@CVPR'20 - Cross-Domain Semantic Segmentation via Domain-Invariant Interactive Relation Transfer
- Manders+@CORR'18 - Simple Domain Adaptation with Class Prediction Uncertainty Alignment
- Munro+@CVPR'20 - Multi-Modal Domain Adaptation for Fine-Grained Action Recognition
- Murez+@CVPR'18 - Image to Image Translation for Domain Adaptation
- Pan+@CVPR'20 - Unsupervised Intra-Domain Adaptation for Semantic Segmentation Through Self-Supervision
- Pei+@AAAI'18 - Multi-Adversarial Domain Adaptation
- Perone+@NEUROIMAGE'19 - Unsupervised domain adaptation for medical imaging segmentation with self-ensembling
- Perrett+@CVPR'19 - DDLSTM: Dual-Domain LSTM for Cross-Dataset Action Recognition
- Piao+@ICCV'19 Semi-Supervised Monocular 3D Face Reconstruction With End-to-End Shape-Preserved Domain Transfer
- Pinhero+@CVPR'18 - Unsupervised Domain Adaptation with Similarity Learning
- Qi+@ICCV'19 - A Novel Unsupervised Camera-Aware Domain Adaptation Framework for Person Re-Identification
- Rozantsev+@PAMI'18 - Beyond Sharing Weights for Deep Domain Adaptation
- Rozantsev+@CVPR'18 - Residual Parameter Transfer for Deep Domain Adaptation
- Roy+@CVPR'19 - Unsupervised Domain Adaptation Using Feature-Whitening and Consensus Loss
- RoyChowdhury+@CVPR'19 - Automatic Adaptation of Object Detectors to New Domains Using Self-Training
- Russo+@ARXIV'17 - From Source to Target and Back: Symmetric Bi-Directional Adaptive GAN

References -5-

- Saito+@ICML'17 - Asymmetric Tri-training for Unsupervised Domain Adaptation
- Saito+@ICLR'18 - Adversarial Dropout Regularization
- Saito+@CVPR'19 - Maximum Classifier Discrepancy for Unsupervised Domain Adaptation
- Saito+@CVPR'19 - Semi-supervised Domain Adaptation via Minimax Entropy
- Sankaranarayanan+@CVPR'18 - Generate To Adapt: Aligning Domains using Generative Adversarial Networks
- Sener+@NIPS'16 - Learning Transferrable Representations for Unsupervised Domain Adaptation
- Shen+@AAAI'18 - Wasserstein Distance Guided Representation Learning for Domain Adaptation
- Shu+@ICLR'18 - A DIRT-T Approach to Unsupervised Domain Adaptation
- Song+@CVPR'19 - Generalizable Person Re-Identification by Domain-Invariant Mapping Network
- Sun+@TASK-CV'16 - Deep CORAL: Correlation Alignment for Deep Domain Adaptation
- Taigman+@ICLR'17 - Unsupervised Cross-Domain Image Generation
- Tzeng+@ARXIV'14 - Deep Domain Confusion: Maximizing for Domain Invariance
- Tzeng+@CVPR'17 - Adversarial Discriminative Domain Adaptation
- Thomas+@ACCV'19 - Artistic Object Recognition by Unsupervised Style Adaptation
- Tobin+@IROS'17 - Domain Randomization for Transferring Deep Neural Networks from Simulation to the Real World
- Tsai+@CVPR'18 - Learning to Adapt Structured Output Space for Semantic Segmentation
- Usman+@ICCV'19 - PuppetGAN: Cross-Domain Image Manipulation by Demonstration
- Vu+@CVPR'19 - ADVENT: Adversarial Entropy Minimization for Domain Adaptation in Semantic Segmentation
- Vu+@ICCV'19 - DADA: Depth-Aware Domain Adaptation in Semantic Segmentation
- Wang+@CVPR'20 - Differential Treatment for Stuff and Things: A Simple Unsupervised Domain Adaptation Method for Semantic Segmentation
- Wei+@CVPR'18 - Person transfer gan to bridge domain gap for person reidentification
- Wu+@ECCV'18 - Dual channel-wise alignment networks for unsupervised scene adaptation,
- Wulfmeier+@IROS'17 - Addressing Appearance Change in Outdoor Robotics with Adversarial Domain Adaptation

References -6-

- Xu+@CVPR'20 - Cross-Domain Detection via Graph-Induced Prototype Alignment
- Xu+@CVPR'20 - Reliable Weighted Optimal Transport for Unsupervised Domain Adaptation
- Yan+@CVPR'17 - Mind the Class Weight Bias: Weighted Maximum Mean Discrepancy for Unsupervised Domain Adaptation
- Yan+@CVPR'20 - Disparity-Aware Domain Adaptation in Stereo Image Restoration
- Yang+@ECCV'18 - Real-to-Virtual Domain Unification for End-to-End Autonomous Driving
- Yang+@CVPR'20 - FDA: Fourier Domain Adaptation for Semantic Segmentation
- Yoo+@ECCV'16 - Pixel-Level Domain Transfer
- Yusuf+@BMVC'18 - Multi-Modal Domain Adaptation for Fine-Grained Action Recognition
- Zellinger+@ICLR'17 - Central Moment Discrepancy (CMD) for Domain-Invariant Representation Learning
- Zhai+@CVPR'20 - AD-Cluster: Augmented Discriminative Clustering for Domain Adaptive Person Re-Identification
- Zhang+@ICCV'17 - Curriculum domain adaptation for semantic segmentation of urban scenes
- Zhang+@CVPR'18 - Domain-Symmetric Networks for Adversarial Domain Adaptation
- Zhang+@CVPR'18 - Collaborative and Adversarial Network for Unsupervised Domain Adaptation
- Zhang+@CVPR'18 - Fully Convolutional Adaptation Networks for Semantic Segmentation
- Zhao+@CVPR'20 - Domain Decluttering: Simplifying Images to Mitigate Synthetic-Real Domain Shift and Improve Depth Estimation
- Zhong+@CVPR'18 - Camera Style Adaptation for Person Re-Identification
- Zhou+@ECCV'18 - Unsupervised Domain Adaptation for 3D Keypoint Estimation via View Consistency
- Zhu+@ICCV'17 - Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks
- Zhu+@ECCV'18 - Penalizing Top Performers: Conservative Loss for Semantic Segmentation Adaptation
- Zhu+@CVPR'19 - Penalizing Top Performers: Conservative Loss for Semantic Segmentation Adaptation
- Zhuang+@IJCAI'15 - Supervised Representation Learning: Transfer Learning with Deep Autoencoders
- Zou+@ECCV'18 - Penalizing Top Performers: Conservative Loss for Semantic Segmentation Adaptation