Institute for Infocomm Research

I²R

REINFORCED CROSS-DOMAIN KNOWLEDGE DISTILLATION ON TIME SERIES DATA

Xu Qing

Institute for Infocomm Research, A*STAR, SG
Nanyang Technological University

Nov-2024

1. Background

- Domain shift widely exist in various TS applications → Unsupervised Domain Adaptation
- However, the power of existing UDA methods heavily depends on the **network capacity**

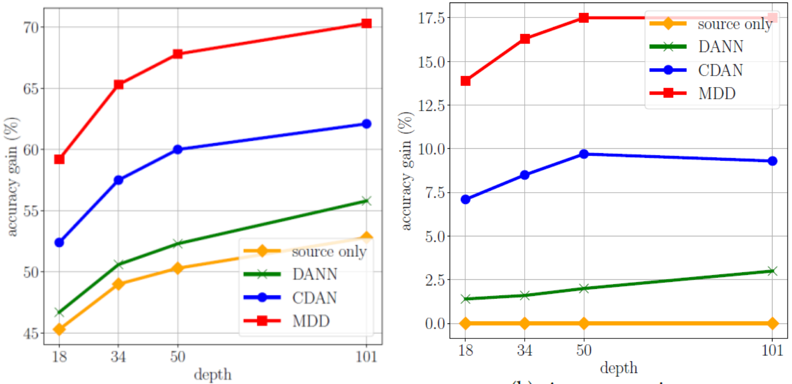

(a) Accuracy (b) Accuracy gain

- Model complexity vs. Constrained Resource

Higher DA performance

↓

Complex network architecture

- Expensive computational and storage cost;
- High power consumption

Deployment ✗

↓

Resource limited devices

* Image from Junguang Jiang, Ximei Wang, Mingsheng Long, and Jianmin Wang. Resource efficient domain adaptation. In Proceedings of the 28th ACM International Conference on Multimedia, pages 2220–2228, 2020.

2. Related Works

- Existing solutions: Integrate Knowledge Distillation into UDA framework
  - Transfer the cross-domain knowledge from Teacher to Student
  - Simultaneously addressing domain shift and model
- Issues of existing solutions:
  - Limited network capacity of student, challenging to capture the fine-grained pattens in data as teacher
  - Teacher’s knowledge on individual target samples may not be always reliable

3. Proposed RCD-KD (Reinforced Cross-Domain Knowledge Distillation)

- A domain discriminator for domain-invariant knowledge transfer via $\mathcal{L}_{DC}$
- A reinforcement learning-based module to learn optimal target sample selection policy for robust knowledge distillation via $\mathcal{L}_{RKD}$
- Two criteria to evaluate teacher's knowledge on target sample: Uncertainty Consistency and Sample transferability

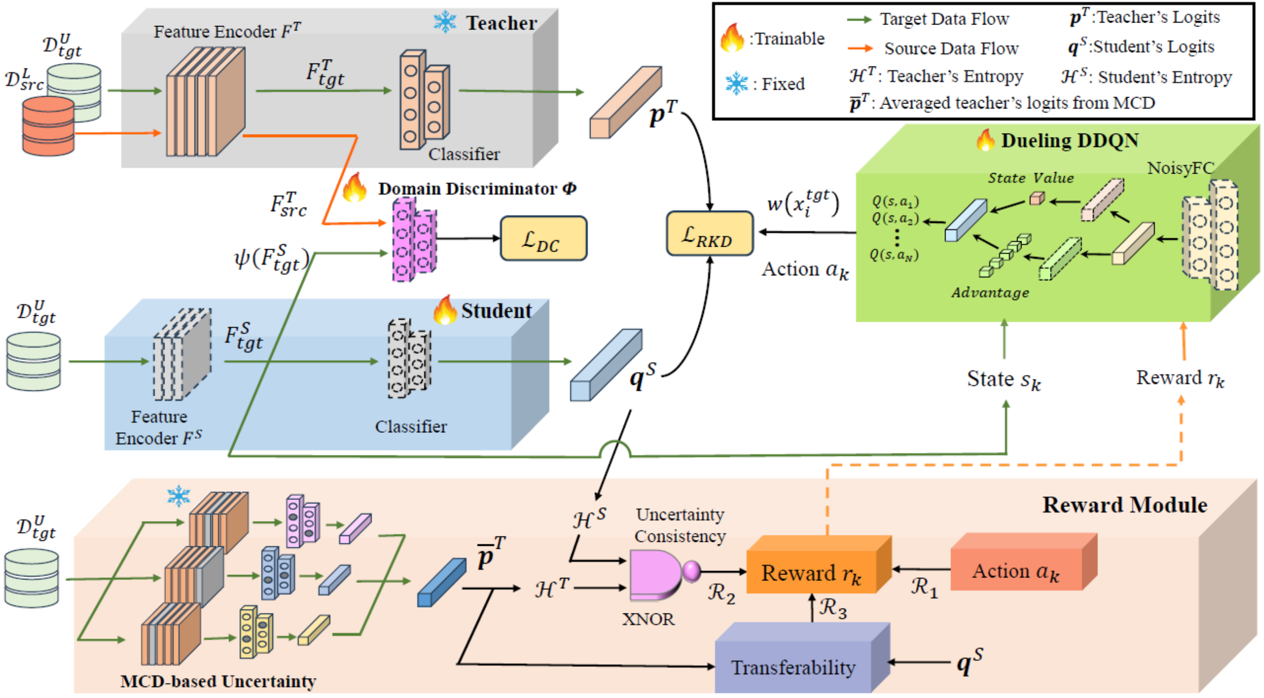

4. Key Components in RL-based Target Sample Selection Module

- **State** $s_k = [F_k^S(x_1^{tgt}), \dots, F_k^S(x_{n_b}^{tgt})]$: Feature representation from student's feature extractor F.

- **Action** $a_k^i \in \{0,1\}$: select or not select the i-th target sample at training step k

- **Reward** $r_k = \alpha_1 * (\mathcal{R}_1 \oplus \mathcal{R}_2 - 0.5) + \alpha_2 * (\mathcal{R}_1 \oplus \mathcal{R}_3 - 0.5)$, where $\oplus$ is exclusive-or
  - Boolean Function $\mathcal{R}_1 = (a_i == 1)$: whether the target sample x_i is retained or not
  - Uncertainty Consistency Reward: whether the student have the same uncertainty level as the teacher for the target sample x_i in a batch. $\odot$ is exclusive-nor operation.

$$\mathcal{R}_2 = (\mathcal{H}_i^S > \frac{1}{n_b} \sum_{j=1}^{n_b} \mathcal{H}_j^S) \odot (\mathcal{H}_i^T > \frac{1}{n_b} \sum_{j=1}^{n_b} \mathcal{H}_j^T)$$

  - Sample Transferability Reward: Whether the target sample x_i is much easier for the student to learn compared to others.

$$\mathcal{R}_3 = (D_i < \frac{1}{n_b} \sum_{j=1}^{n_b} D_j) \ , where\ D_i = KL(\bar{p}_i^T || q_i^S)$$

4. Algorithm for RCD-KD

Algorithm 1 Proposed **RCD-KD**

Input: Teacher T, Student S, adaptation model ψ, domain discriminator Φ, online and target Q-network Q and Q', source data $\mathcal{D}_{src}^{L}$, target data $\mathcal{D}_{tgt}^{U}$ and Replay buffer $\mathcal{M}$

for every epoch **do**
  for every batch $X_{tgt} \in \mathcal{D}_{tgt}^{U}$ and $X_{src} \in \mathcal{D}_{src}^{L}$ **do**
    for episode $k \in [1, K]$ **do**
      Get state s_k, sample action $a_k \sim Q(s_k)$ and update next state s_{k+1}
      Update S and ψ by minimizing $\mathcal{L}$ as Eq. (10)
      Calculate r_k via Eq. (3); Set $d = 0$ if episode end, otherwise $d = 1$
      Store $(s_k, a_k, r_k, s_{k+1}, d)$ to $\mathcal{M}$
      if Φ is not pre-trained **then**
        Fix the parameters in S and ψ and update Φ via minimizing $\mathcal{L}_{adv}$ in Eq. (9)
  Sample a random batch $(s_k, a_k, r_k, s_{k+1}, d)$ from $\mathcal{M}$
  Calculate Q_{est} via Eq. (4) and Q_{tar} via Eq. (5), update Q via Huber loss and Q' via Eq. (6)

$$\mathcal{L} = \mathcal{L}_{DC} + \lambda * \tau^2 * \mathcal{L}_{RKD}.$$

$$\mathcal{L}_{DC} = -\mathbb{E}[log(\Phi(\psi(F^S(x_{tgt}))))].$$

$$\mathcal{L}_{RKD} = \sum_{x \in \mathcal{X}_b} \boxed{w_j} * \sum_i p_i^T log(p_i^T / q_i^S).$$

4. Optimization of DDQN

- Two streams in DDQN: State-value $V(s;\Theta_C,\Theta_V)$ and Advantages $A(s,a;\Theta_C,\Theta_A)$
- Estimated Q-values with online network Q parameterized with $\Theta=\{\Theta_C,\Theta_V,\Theta_A\}$

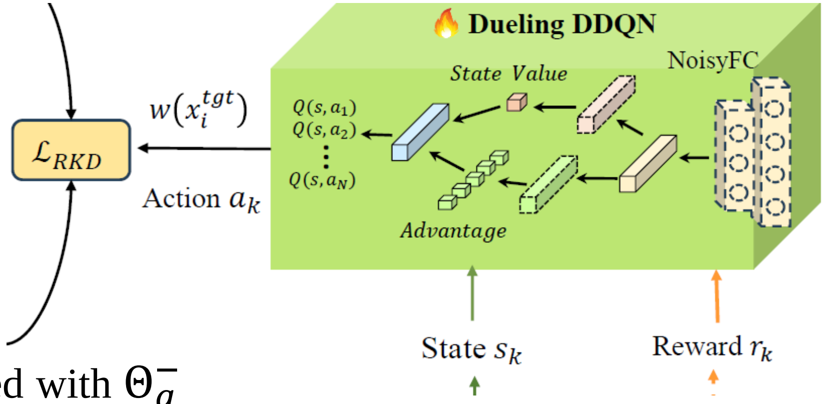

$$Q_{est}=V(s;\Theta_E,\Theta_V)+A(s,a;\Theta_E,\Theta_A)-\frac{1}{2}\sum_{a_i}A(s,a_i;\Theta_E,\Theta_A).$$

- Target Q-values with duplicated network Q' parameterized with Θ_q^-

$$Q_{tar}=r_k+d*\gamma*Q(s_{k+1},\operatorname*{argmax}_{a_{k+1}}Q(s_{k+1},a_{k+1};\Theta);\Theta').$$

- Update Θ_q with the Huber loss between estimated Q-values and target Q-values

- Update Θ' with $\Theta'\leftarrow\delta*\Theta'+(1-\delta)*\Theta.$

5. Experiments - Setup

- Total **4** datasets across **3** different time series tasks.
  - **UCI HAR**: human activity recognition (*i.e.*, walking, walking upstairs, walking downstairs, standing, laying and sitting)
  - **HHAR**: human activity recognition (*i.e.*, biking, sitting, standing, walking, walking upstairs and walking downstairs)
  - **FD**: rolling bearing fault diagnosis, classify bearing health status (*i.e.*, healthy, artificial damages, damages from accelerated lifetime tests)
  - **SSC**: sleep stage classification, identify subject’s sleep stage (*i.e.*, wake, non-rapid eye movement stage, rapid eye movement stage)

- Consider each subject (for HAR and SSC) or operation condition (for FD) as an independent domain, test different transfer scenarios.

6. Benchmark with UDA methods

How well does the student perform when directly applying UDA methods on it?

Table 1: Performance comparison with other UDA methods.

Datasets	Student-Only	Metric-based			Adversarial-based			Ours
		HoMM [6]	MDDA [5]	SASA [36]	DANN [7]	CoDATS [30]	AdvSKM [32]	
HAR	55.94±8.99	83.62±1.82	84.89±6.29	83.37±3.23	82.42±3.82	75.72±8.62	70.72±4.06	**94.68±1.62**
HHAR	58.74±10.79	68.02±6.59	73.26±8.35	77.13±4.09	76.03±1.97	74.64±4.18	63.24±5.99	**82.37±1.84**
FD	66.78±4.38	74.52±6.00	81.80±5.43	86.75±2.39	77.95±8.52	77.54±9.45	77.83±5.71	**92.63±0.62**
SSC	50.39±7.67	59.79±5.51	57.45±3.68	59.36±3.69	57.39±5.51	57.21±5.61	57.28±4.77	**67.49±1.83**

6. Benchmark with KD-based DA methods

Table 2: Marco F1-scores on HAR and HHAR across three independent runs.

Methods	HAR Transfer Scenarios						HHAR Transfer Scenarios					
	$2 \rightarrow 11$	$6 \rightarrow 23$	$7 \rightarrow 13$	$9 \rightarrow 18$	$12 \rightarrow 16$	Avg	$0 \rightarrow 6$	$1 \rightarrow 6$	$2 \rightarrow 7$	$3 \rightarrow 8$	$4 \rightarrow 5$	Avg
Teacher	100.0	100.0	99.92	93.69	81.65	95.05	64.47	94.23	57.22	98.88	97.69	82.50
Student-Only	68.51	59.57	78.88	21.02	51.71	55.94	50.46	65.95	43.22	58.84	75.22	58.74
KD-STDA [11]	98.31	89.55	89.28	67.41	63.13	81.54	46.15	92.19	41.69	96.51	89.79	73.27
KA-MCD [16]	89.46	59.26	63.62	58.93	45.67	63.39	**65.25**	90.59	42.57	85.71	85.48	73.92
MLD-DA [12]	**100.0**	99.11	92.96	82.78	64.08	87.79	61.53	94.32	47.91	91.07	92.74	77.51
REDA [9]	99.44	93.81	92.43	74.55	55.77	83.20	32.05	93.85	36.10	90.24	95.41	69.53
AAD [15]	83.74	90.89	83.05	75.96	61.67	79.06	53.25	81.22	48.35	87.00	86.36	71.24
MobileDA [14]	92.71	90.19	91.39	77.95	64.34	83.32	46.60	93.31	49.13	98.30	96.84	76.84
UNI-KD [4]	**100.0**	96.33	93.20	79.77	64.91	86.84	46.66	**94.89**	**59.20**	**98.45**	**97.42**	79.32
Ours	**100.0**	**100.0**	**99.64**	**92.87**	**80.87**	**94.68**	64.47	94.24	57.59	**98.45**	97.11	**82.37**

Table 3: Marco F1-scores on FD and SSC across three independent runs.

Methods	FD Transfer Scenarios						SSC Transfer Scenarios					
	$0 \rightarrow 1$	$0 \rightarrow 3$	$2 \rightarrow 1$	$1 \rightarrow 2$	$2 \rightarrow 3$	Avg	$0 \rightarrow 11$	$12 \rightarrow 5$	$16 \rightarrow 1$	$7 \rightarrow 18$	$9 \rightarrow 14$	Avg
Teacher	88.36	86.46	88.82	99.84	99.92	92.68	51.43	68.71	73.48	72.48	76.59	68.54
Student-Only	34.94	42.14	75.27	90.41	91.13	66.78	35.62	35.87	60.15	61.24	59.05	50.39
KD-STDA [11]	53.17	50.95	76.76	89.24	98.66	73.76	43.75	53.45	49.04	67.23	65.56	55.81
KA-MCD [16]	57.96	65.26	61.66	81.75	91.79	71.68	**50.85**	56.73	51.01	64.18	65.95	57.74
MLD-DA [12]	78.16	75.49	83.34	99.86	96.83	86.74	45.36	66.17	58.37	63.87	70.71	60.90
REDA [9]	86.70	81.08	88.98	92.35	88.77	87.58	44.07	52.01	60.14	60.46	64.67	56.27
AAD [15]	52.50	60.00	80.86	89.84	95.99	75.84	32.71	62.92	63.34	64.46	72.15	59.12
MobileDA [14]	76.19	58.77	83.74	97.56	97.84	82.82	41.83	57.14	59.41	64.38	61.55	56.86
UNI-KD [4]	78.85	82.68	**92.14**	97.29	**99.34**	90.06	44.48	60.13	62.99	71.03	72.21	62.17
Ours	**89.88**	**85.63**	88.57	**99.92**	99.13	**92.63**	49.73	**70.74**	**72.14**	**71.73**	**76.95**	**68.26**

7. Ablation Study: Different T-S pairs and Teacher Generation

- Different T-S pairs

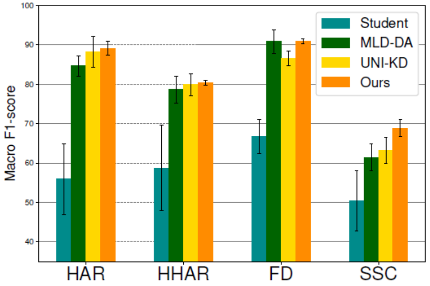

Figure 2: **TCN →1D-CNN**.

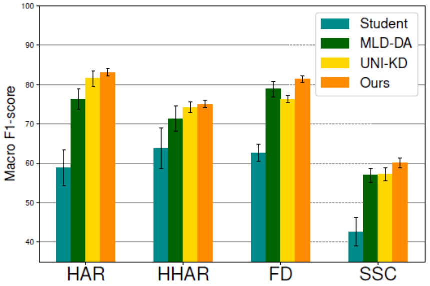

Figure 3: **Resnet-34 → Resnet-18**.

- Different Teacher Generation

Table 4: Teacher with different UDA methods.

T-Types	HAR	HHAR	FD	SSC
MDDA	89.16	81.65	90.45	61.62
SASA	88.16	80.39	90.83	63.25
CoDATS	93.39	**82.98**	91.08	66.89
DANN (ours)	**94.68**	82.37	**92.63**	**68.26**

7. Ablation Study: Reinforced Sample Selection Ablation

Table 6: Reinforced sample selection ablation. “Full samples" denotes utilizing whole target samples for KD; ‘$\mathcal{R}_2$’, ‘$\mathcal{R}_3$’ denote directly utilizing proposed uncertainty and transferability for sample selection; ‘$\mathcal{R}_2^\dagger$’, ‘$\mathcal{R}_3^\dagger$’ denote utilizing RL with $\mathcal{R}_2$ and $\mathcal{R}_3$ as reward for sample selection; $(\mathcal{R}_2 + \mathcal{R}_3)^\dagger$ is **our** proposed method.

Datasets	Full Samples	Partial Samples					
		$\mathcal{R}_2$	$\mathcal{R}_2^\dagger$	$\mathcal{R}_3$	$\mathcal{R}_3^\dagger$	$\mathcal{R}_2 + \mathcal{R}_3$	$(\mathcal{R}_2 + \mathcal{R}_3)^\dagger$
HAR	89.32	91.65	93.91	92.31	93.96	93.53	**94.68**
HHAR	78.99	78.30	81.73	80.33	82.29	81.04	**82.37**
FD	89.13	90.17	91.93	89.51	91.08	91.85	**92.63**
SSC	60.65	63.16	62.98	65.81	67.49	65.20	**68.26**

7. Ablation Study: Computational Complexity

Table 7: Comparison of Computational Complexity.

Methods	KD-STDA	KA-MCD	MLD-DA	REDA	AAD	MobileDA	UNI-KD	**Ours**
Time (sec)	1.68	4.55	1.91	1.78	0.91	1.28	3.26	16.42

Thanks for your time!