

Progressive Domain Adaptation under Sensor Drift for Gas Classification with Limited Calibration Labels

Han Fan^{1,2}, Achim J. Lilienthal^{1,2}

¹ Chair of Perception for Intelligent Systems (PerInS), Munich Institute of Robotics and Machine Intelligence (MIRMI), Technical University of Munich (TUM), Germany

² Robot Navigation and Perception Lab (RNP), Örebro University, Sweden
han.fan@tum.de, achim.j.lilienthal@tum.de

Summary: Metal-oxide electronic noses suffer from sensor drift, in which sensor response characteristics change over time. This creates a data distribution shift (domain gap) that severely degrades gas classification accuracy. To address this drift effect, domain adaptation techniques reduce the distribution discrepancy between the initial source domain and drift-affected target measurements. We study batch-wise adaptation across temporally ordered drift batches, where calibration labels are extremely scarce. To ensure robust adaptation under unpredictable drift patterns, we train a diverse pool of conditional domain-adversarial experts. We then fuse their predictions using a confidence-aware Expectation-Maximization procedure. This fusion mechanism acts as a smart consensus, estimating each expert's reliability using the limited calibration labels without requiring additional annotations. On the UCI drift benchmark, this budget-respecting approach consistently improves target accuracy over a tuned Drift Correction Autoencoder (DCAE) baseline, particularly on later, more severely drifted batches.

Keywords: electronic nose; sensor drift; transfer learning; domain adaptation.

Introduction

Metal-oxide (MOX) electronic noses (e-noses) enable low-cost gas identification, but their long-term reliability is limited by sensor drift and instrumental variation. Such effects induce distribution shifts between an initial training condition and later operating conditions, leading to substantial accuracy degradation in gas classification models. Existing drift-compensation strategies rely on different assumptions about the availability of the target domain. Correction-centric transfer learning can exploit paired transfer measurements to learn drift-corrected yet discriminative representations (e.g., drift correction autoencoder (DCAE)) [1]. At the other extreme, target-domain-free methods (e.g., TDACNN) avoid using any target-domain samples during training and prediction, which improves deployability but prevents adaptation to batch-specific drift patterns as conditions evolve.

In practice, continuously deployed e-noses yield abundant unlabeled measurements, but physical recalibration is costly, restricting us to a strict budget of calibration pairs. Conditional Adversarial Domain Adaptation (CDAN) leverages unlabeled data to align source and target distributions, conditioned on class predictions [2]. However, when drift is severe and heterogeneous, aggressively enforcing domain alignment may degrade target-domain classification performance and lead to negative transfer [3].

To achieve robust gas classification under a strict calibration budget, we propose a two-

stage framework. First, we train a diverse pool of CDAN experts. By intentionally varying their adaptation components, we induce a variety of drift-compensation behaviors. Second, we fuse their predictions using a confidence-aware Dawid–Skene EM procedure. This semi-supervised procedure estimates confidence-specific expert confusion matrices from the limited calibration labels and uses them to infer a reliability-aware fused class posterior.

Method

We treat the gas sensor drift as a statistical domain shift between the initial calibration data (source domain, $\mathcal{S}$) and the incoming batch (target domain, $\mathcal{U}_b$). Therefore, by applying domain adaptation to align these two distributions, we can effectively reduce the drift effect. We propose an approach that combines conditional domain-adversarial learning with ensemble learning. To handle heterogeneous drift, we leverage a multi-expert pool to capture diverse adaptation trajectories, then employ a probabilistic consensus mechanism to resolve expert disagreements using the limited calibration data. Our approach consists of two stages:

A. Mechanistically Diverse Domain-Adversarial Experts

Because drift patterns are unpredictable, a single adaptation strategy may over-correct or under-correct the distribution shift. To mitigate this, we train a pool of experts $m \in \{1, \dots, M\}$

using the following loss function:

$$\mathcal{L}^{(m)} = \mathcal{L}_{\text{cls}} + \lambda_{\text{adv}}^{(m)} \mathcal{L}_{\text{CDAN}} + \lambda_{\text{pair}}^{(m)} \mathcal{L}_{\text{pair}} + \lambda_{\text{pst}}^{(m)} \mathcal{L}_{\text{pst}} \quad (1)$$

Here, $\mathcal{L}_{\text{cls}}$ ensures basic classification on source data; $\mathcal{L}_{\text{CDAN}}$ encourages conditional alignment between source and target representations [2]; $\mathcal{L}_{\text{pair}}$ aligns the few paired source–target transfer samples; and $\mathcal{L}_{\text{pst}}$ performs self-training with predicted high-confidence pseudo-labels. To cover a wide spectrum of drift-compensation behaviours, we systematically vary the hyperparameters (λ) and conditioning strategies across the experts.

B. Confidence-aware Fusion

Since different experts may excel under different drift conditions, we use a semi-supervised Dawid–Skene Expectation-Maximization (EM) procedure to dynamically find the best consensus. For every unlabeled testing measurement, the experts provide class-probability distributions. ConfDS estimates confidence-specific confusion matrices and iteratively infers a fused posterior over gas classes using Dawid–Skene expectation–maximization. To allow an expert’s reliability to differ between confident and uncertain predictions, we separate each expert’s predictions into high-confidence and low-confidence bins (using a threshold τ), evaluating their reliability separately. By anchoring this estimation on the tiny budget of true calibration labels, the EM algorithm updates the reliability weights and outputs a robust, fused final classification.

Results and Conclusion

We use the open-source UCI Gas Sensor Array Drift Dataset [4] for evaluation. We use the first batch as the source domain and perform batch-wise adaptation and evaluation on the temporally ordered target batches b02–b10. Only 0.62% target measurements are labelled for the confidence-aware fusion (the per-batch fraction varies with batch size). We report target classification accuracy and its average over target batches. The result is obtained by repeating the full pipeline over 6 runs (3 training seeds $\times$ 2 pair-shuffle seeds).

Fig. 1 summarizes the per-batch classification accuracy. Overall, the proposed CDAN diverse-expert pool with ConfDS fusion achieves a higher mean target accuracy than the DCAE baseline trained under the same budget, with the advantage becoming most pronounced on later/harder batches, namely b06–b09. Both our work and the DCAE outperform the target-domain-free baseline implemented by TDACNN [5], which illustrates the typical accuracy degradation under sensor drift when no target-domain




Figure 1: Per-batch classification accuracy on the UCI drift dataset. Our ConfDS fusion (0.899 ± 0.009 overall) consistently outperforms the DCAE baseline (0.885 ± 0.008) under the same strict calibration budget and the target-domain-free baseline, with the largest gains observed on b06–b09.

adaptation is performed.

In conclusion, our results show that combining diverse conditional domain-adaptation experts with confidence-aware fusion can improve gas classification robustness under temporally varying sensor drift and limited target supervision. For reproducibility, the implementation of the proposed framework is publicly available at <https://github.com/TUM-PercInS/confds-enose-domain-adaptation.git>. As a next step, we will validate the approach on long-term robotic olfaction deployments with operationally realistic calibration procedures.

References

- [1] Ke Yan and David Zhang. Correcting instrumental variation and time-varying drift: A transfer learning approach with autoencoders. *IEEE Transactions on Instrumentation and Measurement*, 65(9):2012–2022, 2016.
- [2] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I Jordan. Conditional adversarial domain adaptation. *Advances in neural information processing systems*, 31, 2018.
- [3] Han Zhao, Remi Tachet Des Combes, Kun Zhang, and Geoffrey Gordon. On learning invariant representations for domain adaptation. In *International conference on machine learning*, pages 7523–7532. PMLR, 2019.
- [4] Alexander Vergara, Shankar Vembu, Tuba Ayhan, Margaret A Ryan, Margie L Homer, and Ramón Huerta. Chemical gas sensor drift compensation using classifier ensembles. *Sensors and Actuators B: Chemical*, 166:320–329, 2012.
- [5] Yuelin Zhang, Sihao Xiang, Zehuan Wang, Xiaoyan Peng, Yutong Tian, Shukai Duan, and Jia Yan. Tdaccnn: Target-domain-free domain adaptation convolutional neural network for drift compensation in gas sensors. *Sensors and Actuators B: Chemical*, 361:131739, 2022.