



Relative Norm Alignment for Tackling Domain Shift in Deep Multi-modal Classification

Mirco Planamente^{1,2,3} · Chiara Plizzari¹ · Simone Alberto Peirone¹ · Barbara Caputo^{1,3} · Andrea Bottino¹

Received: 3 April 2023 / Accepted: 5 January 2024
© The Author(s) 2024

Abstract

Multi-modal learning has gained significant attention due to its ability to enhance machine learning algorithms. However, it brings challenges related to modality heterogeneity and domain shift. In this work, we address these challenges by proposing a new approach called Relative Norm Alignment (RNA) loss. RNA loss exploits the observation that variations in marginal distributions between modalities manifest as discrepancies in their mean feature norms, and rebalances feature norms across domains, modalities, and classes. This rebalancing improves the accuracy of models on test data from unseen (“target”) distributions. In the context of Unsupervised Domain Adaptation (UDA), we use unlabeled target data to enhance feature transferability. We achieve this by combining RNA loss with an adversarial domain loss and an Information Maximization term that regularizes predictions on target data. We present a comprehensive analysis and ablation of our method for both Domain Generalization and UDA settings, testing our approach on different modalities for tasks such as first and third person action recognition, object recognition, and fatigue detection. Experimental results show that our approach achieves competitive or state-of-the-art performance on the proposed benchmarks, showing the versatility and effectiveness of our method in a wide range of applications.

Keywords Multi-modal learning · Norm alignment · Domain generalization · Unsupervised domain adaptation

1 Introduction

Humans have the ability to perceive the world around them through signals that come from multiple sensory systems. Our perceptual experiences can be visual, auditory, tac-

tile, olfactory, and gustatory. Psychologists and neurologists agree that our perception does not depend on a single modality at a single time, but is fundamentally multi-modal in nature (Bertelson & Gelder, 2004; Stein et al., 2002). Moreover, the interpretation of data from one sensory channel is influenced by data from other modalities (Driver & Spence, Jul. 1998; O’Callaghan, 2012).

The same ability to effectively process and integrate information from multiple sensory channels has been shown to significantly improve the performance of current machine learning algorithms. For example, recent video understanding models use complementary audio-visual (Zhu et al., 2021) and appearance-motion information (Munro & Damen, 2020; Sevilla-Lara et al., 2019; Sun et al., 2018) to improve accuracy and generalization performance. Object recognition algorithms use depth information to extract more accurate information and classify objects more effectively (Loghmani et al., 2020; Bo et al., 2011), and so on.

Despite its potential benefits, Multi-Modal Learning (MML) also comes with some challenges, such as learning how to summarize data while retaining their complemen-

Communicated by Bodo Rosenhahn.

✉ Mirco Planamente
mirco.planamente@polito.it

Chiara Plizzari
chiara.plizzari@polito.it

Simone Alberto Peirone
simone.peirone@polito.it

Barbara Caputo
barbara.caputo@polito.it

Andrea Bottino
andrea.bottino@polito.it

¹ DAUIN, Politecnico di Torino, Torino, Italy

² Italian Institute of Technology, Genoa, Italy

³ Consortium Cini, Rome, Italy

tary information (Wang et al., 2020) or understanding how to effectively combine information from multiple modalities when making a prediction (Baltrušaitis et al., 2019). The same issue is addressed in Yang et al. (2022) when data from multiple views are used. Heterogeneity between modalities is another critical issue, as the difference between their marginal distributions may prevent the model from learning equally from all of them (Razzaghi et al., 2021). Another well-known problem in the literature is the so-called “domain shift”, i.e., that a model trained on a labeled source dataset does not generalize well to an unseen target dataset. Different modalities may be affected differently by the domain shift (Lv et al., 2021). For example, when using audio-visual data for egocentric action recognition, the action “cut” in a cooking scenario may reveal differences between domains (Planamente et al., 2022), as cutting boards in different kitchens may differ in their visual and auditory impressions (e.g., wooden cutting board vs. plastic cutting board), different types of food may be cut, and so on. This highlights the need for robust models that can handle variation across modalities and domains.

To address both the cross-modal and cross-domain challenges in MML, we recently proposed in Planamente et al. (2022) a simple but effective audio-visual approach in the context of egocentric action recognition. We observed that differences in the marginal distributions of the audio and visual modalities could lead to variations in feature informativeness that do not only negatively affect the training process and lead to suboptimal performance, but also typically translate into discrepancies between the mean norms of their features. This imbalance in norms leads the network to “favor” the modality with the larger features, which prevents the model from fully exploiting the synergies and complementarities between modalities and reduces its generalization capabilities (Barbato et al., 2021).

To tackle this issue, in Planamente et al. (2022) we proposed to reduce such imbalance with a simple loss called *Relative Norm Alignment* (RNA) loss. In the Domain Generalization (DG) setting, i.e., when the model does not have access to the target data at training time, this loss attempts to align the average norms of the different modalities to a common value. This objective also leads to successful transfer between source and target (Zhou et al., 2020; Zheng et al., 2018; Xu et al., 2019; Barbato et al., 2021). In the Unsupervised Domain Adaptation (UDA) setting, i.e., when target data are available during training, RNA is defined as the sum of two domain-specific terms that aim to achieve a cross-modality norm balance on both source and target domains. However, in this setting, the RNA loss operated separately on the two domains (*cross-modal alignment*, Fig. 1), resulting in discrepancies between the mean feature norms of the two. This discrepancy can be explained by the presence of

domain-specific features from the source domain that may have low activations in the target domain.

To improve the effectiveness of RNA, in this work we extend it to independently align feature norms for each modality across domains (*cross-modal alignment*, Fig. 1) so that the network can prioritize more transferable features (Xu et al., 2019). In addition, we address the problem of imbalanced feature norms between classes by introducing an intra- and inter-domain alignment component per class (*per-class alignment*, Fig. 1), resulting in improved overall accuracy.

Furthermore, we combine RNA with two additional components in UDA settings. First, we incorporate an adversarial loss to improve domain-invariant feature learning. Second, we observe that the original RNA loss only affects the modality embedding models and neglects the classification layers. To mitigate the prediction uncertainty in the target domain, we extend the training loss of the model with an Information Maximization term that uses pseudo-labels on target data.

The solution we propose differs from previous approaches in that it is simple, it does not require changes to the training process and, differently from recent contrastive-learning based approaches (Sahoo et al., 2021; Song et al., 2021; Kim et al., 2021), it does not require effective mining of hard negative samples. This makes our solution a desirable choice for a broader range of modalities and tasks. In particular, we extend the audio-visual loss proposed in Planamente et al. (2022) to a variety of visual and nonvisual modalities (optical flow, event data, depth, EGG, facial keypoints) and to a variety of tasks, including first- and third-person action recognition, object recognition, and fatigue detection. Despite its simplicity, experiments show that our approach performs equally well, if not better, than existing methods, with a leaner and more efficient implementation.

In summary, the main contributions of this work are as follows:

- it updates the definition of RNA to improve the transferability of features between domains in DG and UDA settings;
- it introduces the use of pseudo-labeling to regularize predictions in the context of transfer learning between source and target;
- it addresses the challenges of multi-modal domain shift by extending our analysis to multiple modalities and multiple tasks;
- it presents a comprehensive analysis and ablation of our approach in both DG and UDA settings, showing state-of-the-art or competitive performances on all benchmarks.

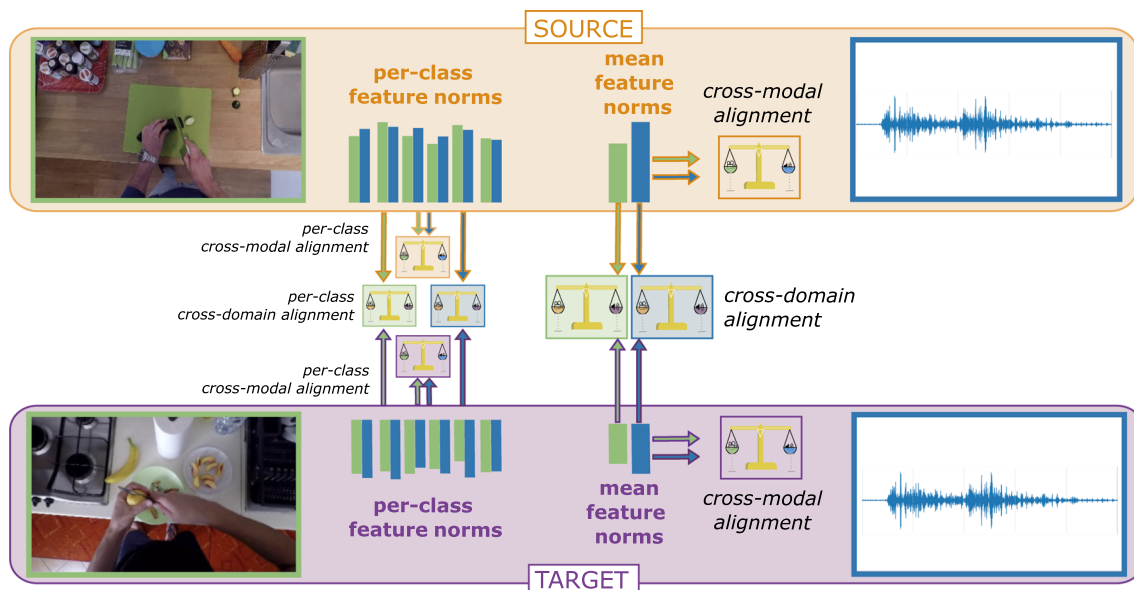


Fig. 1 Overview of Relative Norm Alignment (RNA) loss for RGB and audio modalities. Given visual and audio input from both source and target domains, we perform an alignment at feature level by re-balancing (i) the mean feature norms of visual and audio modalities (*cross-modal*

alignment, \mathcal{L}_{RNA}^s), (ii) per-class mean feature norms of visual and audio modalities (*per-class alignment*, \mathcal{L}_{RNA}^c) and (iii) mean feature norms of source and target features independently for each modality (*cross-domain alignment*, \mathcal{L}_{RNA}^{mod})

2 Related Work

MML has gained popularity due to its potential for better performance, robustness, and deeper understanding. Previous surveys Baltrušaitis et al. (2019), Rahate et al. (2022) have discussed the challenges and opportunities of MML. Here we focus on the problem of generalizing MML across domains, and in particular explore computer vision applications that are most relevant to the experiments in our work.

2.1 Domain Adaptation

In DG, the goal is to build a model that uses knowledge from one or multiple source domains, without having access to data from the target domain during training, to improve the generalization performance of the model to any unseen domain. In such a setting, the lack of knowledge about the target distributions prevents the possibility to estimate the domain discrepancy between source and target domains. Computer vision based DG approaches have mainly focused on image data and can be broadly classified into several categories. Feature-based methods aim to learn domain-invariant representations by aligning domain distributions with metrics such as MMD (Li et al., 2018; Gretton et al., 2012) or CORAL (Sun & Saenko, 2016), or with domain adversarial networks (Ganin et al., 2016). Data-based methods increase the amount of training data to prevent overfitting or use style transfer to reduce the domain sensitivity (Volpi et al., 2018; Zhang et al., 2022; Chen et al., 2022a, b; Wang et al., 2020; Xu et al.,

2020). Meta-Learning methods simulate the shift in distributions between domains (Balaji et al., 2018; Dou et al., 2019; Li et al., 2019). Self-Supervision (Bucci et al. 2021) uses auxiliary tasks to learn generalizable representations.

In the context of video data, VideoDG (Yao et al. 2021) observes that it is important to find a balance between the ability to generalize and the ability to discriminate. To achieve this, the relationships between frames in the source domain are extended to ensure that they can generalize to potential target domains while maintaining their discriminative capabilities.

As for UDA methods (which can benefit from unlabeled target data available during training), they can be broadly divided into two categories: discrepancy-based and adversarial-based methods. Discrepancy-based methods minimize a distance metric between the source and target distributions (Xu et al., 2019; Saito et al., 2018; Long et al., 2015). Adversarial-based methods, on the other hand, use adversarial training to align source and target distributions (Deng et al., 2019; Tang & Jia, 2020). Another research direction focuses on incorporating self-supervised learning as an auxiliary task to improve feature learning, as in Bucci et al. (2021).

While the aforementioned approaches have mainly been applied to standard image classification tasks, there has also been a significant amount of research on UDA for video-related tasks, such as action detection (Agarwal et al., 2020), segmentation (Chen et al., 2020), and classification (Chen et

al., 2019; Munro & Damen, 2020; Choi et al., 2020a; Jamal et al., 2018; Pan et al., 2020; Song et al., 2021).

In video classification, several methods have been proposed to align the temporal dynamics of the feature space. TA³N (Chen et al. (2019)) uses a multi-level adversarial framework with temporal relation and attention mechanisms to achieve this goal. TCoN (Pan et al. 2020) aligns feature distributions between source and target domains with a cross-domain co-attention mechanism that focuses on aligning temporal relationship features to increase robustness across domains. In Choi et al. (2020a), the network is trained to solve an auxiliary self-supervised task on source and target data. SAVA (Choi et al. 2020b) addresses the domain adaptation problem by proposing to use clip order prediction as an auxiliary task to be solved in both source and target domains. In addition, Contrastive Learning (CL) methods have also been proposed for UDA in video analysis. For example, CoMix (Sahoo et al. 2021) introduced a new framework for contrastive learning that aims to learn discriminative invariant feature representations.

2.2 Multi-modal Adaptation and Generalization

Several methods have been proposed to exploit the availability of multiple modalities for domain adaptation. They can be divided into three main categories: adversarial approaches, co-training, and contrastive learning based methods.

Adversarial-based approaches, such as MDANN (Qi et al., 2018) and AUDA (Liu et al., 2021), focus on learning discriminative and domain adaptive features under an adversarial objective, showing their effectiveness in cross-domain emotion recognition using audio-visual data and cross-media retrieval using images and text from different domains. MM-SADA (Munro and Damen 2020) is another approach that extends adversarial alignment to a self-supervised task based on modality correspondence.

Co-training methods such as DLMM (Lv et al., 2021) and XM-UDA (Jaritz et al., 2020) exploit the diverse properties of the different modalities by treating the classifiers of the various modalities as a set of teacher/student models trained with a curriculum learning approach. These methods have been applied to tasks such as event recognition using audio-visual data, fatigue detection using EEG signals and facial keypoints, and action recognition using RGB images and optical flow.

Contrastive learning based methods such as STCDA (Song et al., 2021) and the approach described in Kim et al. (2021) exploit the complementarity of different modalities to regularize both cross-modal and cross-domain feature representations. They treat each modality as a view and perform contrastive learning across modalities and domains to align representations between source and target domains in each modality. CIA (Yang et al., 2022) uses cross-modal

interaction and generative modelling to align cross-domain representations.

RNA-Net (Planamente et al., 2022) addresses multi-modal video DG by using both audio and RGB features, but recognizes that the simple fusion of multi-modal information may not improve generalizability. To overcome this problem, a cross-modal audio-visual Relative Norm Alignment (RNA) loss is proposed to align the relative feature norms of audio and visual modalities from source domains, resulting in domain-invariant audio-visual features. In this work, we further extend this approach to improve feature transferability across domains in both UDA and DG settings, and address the issues of multi-modal domain shift across different tasks and datasets.

2.3 Norm Alignment

Several works highlighted the existence of a strong correlation between the mean feature norms and the amount of “valuable” information for classification (Zheng et al., 2018; Wang et al., 2017; Ranjan et al., 2017) and the negative impact of different feature norms on multiview clustering approaches (Peng et al., 2019). In particular, the cross-entropy loss has been shown to promote well-separated features with a high norm value (Wang et al., 2017). Starting from this observation, the authors of Xu et al. (2019) show that the main reason behind performance degradation on unseen data is the reduction in feature norms compared to the source domain. This stems from the fact that the supervision on the source domain causes the classifier to rely on domain-specific features that may not be present in the target domain, thus reducing the activations in the representation of the target features and consequently the norms of the target features. To address this problem, Xu et al. (2019) introduced a loss that forces the norms between the two domains to adapt to increasingly larger scalars, resulting in improved transfer between domains.

Similarly, Barbato et al. (2021) proposed a regularization objective that promotes uniform feature norms between source and target representations while also inducing progressively higher norm values. Furthermore, they introduced an inter-class norm alignment objective, based on the observation that classes with higher confidence are associated with larger feature norms, to soften distribution biases towards the most frequent classes, whose higher classification confidence is typically associated with larger feature norms.

Subsequent works have demonstrated the effectiveness of incorporating this regularization term into various approaches to learn domain-invariant features, such as the adversarial distribution adaptation network proposed in Zhou et al. (2020) and the hierarchical transfer network described in Yang et al. (2021).

In this work, we apply the concept of norm alignment to domain adaptation by extending it to a multi-modal setting, where the alignment is performed not only between domains, but also across modalities and classes. This allows us to better handle the complexity of multi-modal data and improve the transferability of features across different domains and modalities.

3 Proposed Method

In the following, we detail the proposed Relative Norm Alignment (RNA) loss, which aims to mitigate the domain shift in MML by aligning the mean feature norms from different modalities (*cross-modal alignment*) and from different domains (*cross-domain alignment*), both globally and at class level.

3.1 Intuition and Motivation

Joint training of multi-modal models may result in sub-optimal synergies between the different modalities. This observation has been theoretically demonstrated in Huang et al. (2022), showing that naive joint training prevents efficient learning from all modalities. From an optimization perspective, the modality with better performance contributes to lower joint discriminative loss and dominates the training progress, while smaller gradient updates are propagated through the other modalities, leading to an under-optimized situation in which the dominant modality learn faster than the others (Peng et al., 2022). In turn, the cross entropy loss encourages the network to learn more separable features (Xu et al., 2019), thus increasing their feature norms unevenly. This problem becomes particularly relevant in cross-domain scenarios, where the accuracy drop is further exacerbated by domain shift.

For this reason, in this work we introduce a new loss function based on the mean features norm of the different modalities. This loss promotes balanced learning and synergistic integration of modalities. By addressing the issues of modality imbalances and domain shift, RNA improves the model’s ability to effectively exploit multi-modal information and improve overall performance.

3.2 Setting

Suppose we observe data $\mathcal{X}_S = \{(x_{s,i}, y_{s,i})\}_{i=1}^{n_s}$ from a source distribution \mathcal{S} , where n_s is the total number of samples, each associated with a label $y_{s,i}$ from the label space \mathcal{Y}_S . Each sample $x_{s,i}$ contains multiple modalities, i.e., $x_{s,i} = \{x_{s,i}^1, \dots, x_{s,i}^M\}$, where $x_{s,i}^m$ denotes the m^{th} modality of the i^{th} sample and M is the number of modalities. The target domain \mathcal{T} comprises n_t annotated target samples $\mathcal{X}_T = \{x_{t,i}\}_{i=1}^{n_t}$,

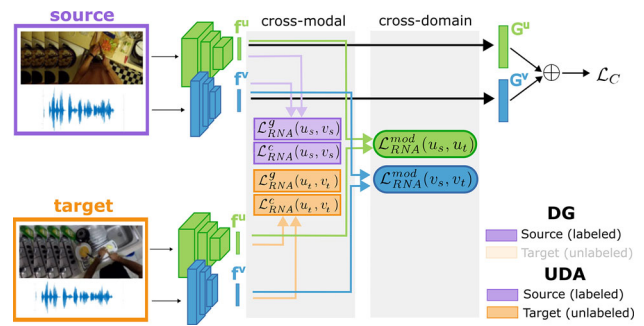


Fig. 2 Labeled source and unlabeled target samples from the modalities u (e.g., visual) and v (e.g., audio) are fed to the respective feature extractors. \mathcal{L}_{RNA} aims to balance the relative feature norms of the two modalities, through a combination of the (domain-specific) cross-modal components (\mathcal{L}_{RNA}^g and \mathcal{L}_{RNA}^c) and the cross-domain ones (\mathcal{L}_{RNA}^{mod}) in each u and v modality. In DG, only the components computed on the source are used

each characterized by the same M modalities of the source samples (i.e., $x_{t,i} = \{x_{t,i}^1, \dots, x_{t,i}^M\}$).

We assume that the distributions of all involved domains are different, i.e., $\mathcal{D}_{d_1}^j \neq \mathcal{D}_{d_2}^k$, where d_1 and d_2 are the domains (source or target) and j and k represent different modalities on the same domain and the same or different modalities on different domains. We also assume that the label space is shared between sources and targets, i.e., $\mathcal{Y}_S = \mathcal{Y}_T$.

3.3 RNA for Domain Adaptation

In the following, without loss of generality, we consider a single-source single-target problem in which two modalities are available. In Sect. 3.5 we show how the approach can be extended to any number of modalities.

We denote each input sample i as $x_i = (x_i^u, x_i^v)$, where u and v represent the two modalities (e.g., visual and audio modality). As shown in Fig. 2, each input modality m is fed to a separate features extractor F^m . The features $f_i^m = F^m(x_i^m)$ are then processed by a classifier G^m , which outputs the score predictions for the m^{th} modality of the i^{th} sample. Finally, the prediction scores from all modalities are combined using a *late fusion* approach to obtain the final predictions. In UDA settings, the F^m feature extractors are shared between source and target.

As previously mentioned, in this work we extend the approach introduced in Planamente et al. (2022), which proposes to train the entire architecture by minimizing the following loss:

$$\mathcal{L} = \mathcal{L}_C + \mathcal{L}_{RNA}^g$$

where \mathcal{L}_C is the standard *cross-entropy loss* on source data. The latter aims at *globally* minimizing the difference between

the feature norms of the two modalities and is defined as:

$$\mathcal{L}_{RNA}^g(u, v) = \lambda_g \left(\frac{\mathbb{E}[h(X^u)]}{\mathbb{E}[h(X^v)]} - \frac{\mathbb{E}[h(X^v)]}{\mathbb{E}[h(X^u)]} \right)^2 \quad (1)$$

where $h(x_i^m) = (\|\cdot\|_2 \circ F^m)(x_i^m)$ is the L_2 -norm of m^{th} modality features of the i^{th} sample, $\mathbb{E}[h(X^m)] = 1/B \sum_{x_i^m \in \mathcal{X}^m} h(x_i^m)$ is the average norm for the m^{th} modality of the B samples composing the batch, and λ_g weights \mathcal{L}_{RNA}^g . To ensure that all features have the same dimension, we project them to a common shape using a fully connected layer when this condition is not met.

In DG, the RNA objective is defined as $\mathcal{L}_{RNA} = \mathcal{L}_{RNA}^g(\mathcal{S})$ while in UDA $\mathcal{L}_{RNA} = \mathcal{L}_{RNA}^g(\mathcal{S}) + \mathcal{L}_{RNA}^g(\mathcal{T})$, where $\mathcal{L}_{RNA}^g(\mathcal{S})$ and $\mathcal{L}_{RNA}^g(\mathcal{T})$ are the loss in Eq. 1 applied to the source and target domains, respectively.

The dividend/divisor structure of \mathcal{L}_{RNA}^g promotes a relative adjustment between the global norm of the two modalities aimed at achieving an *optimal equilibrium* between the two. The square of the difference forces the network to take larger steps when the ratio of the two modality norms is too different, leading to faster convergence. We note that Eq. 1 redefines the loss presented in Planamente et al. (2022) to ensure a symmetric form (i.e., $\mathcal{L}_{RNA}^g(u, v) = \mathcal{L}_{RNA}^g(v, u)$).

3.4 RNA Extensions

While the results in Planamente et al. (2022) show the effectiveness of \mathcal{L}_{RNA}^g in reducing domain shift, the formulation in Eq. 1 has two major limitations. First, the *global* cross-modal alignment performed by \mathcal{L}_{RNA}^g may also lead to unbalanced norms between modalities at the class level, which in turn tends to favor one modality over the others when making decisions about particular classes. Second, in UDA, the alignment is performed separately for each domain. As a result, the average feature norms may still show large differences between source and target domains. These differences can be attributed to the presence of domain-specific features that originate from training in the source domain and may have low activations in the target domain (Xu et al., 2019; Barbato et al., 2021), affecting overall accuracy.

To address both problems, we propose the following extensions to the RNA formulation. First, we introduce an intra-domain class constraint \mathcal{L}_{RNA}^c to address the cross-modal norm imbalance at class level, defined as follows:

$$\mathcal{L}_{RNA}^c(u, v) = \lambda_c \sum_{c=1}^C \left(\frac{\mathbb{E}[h(X_c^u)]}{\mathbb{E}[h(X_c^v)]} - \frac{\mathbb{E}[h(X_c^v)]}{\mathbb{E}[h(X_c^u)]} \right)^2$$

where λ_c weights the loss, and $\mathbb{E}[h(X_c^m)]$ denotes the average norm of the features of modality m for samples of class c , with C the total number of classes. We note that in computing

\mathcal{L}_{RNA}^c for the target, the pseudo-labels are used to assign the target samples to classes.

The second extension of \mathcal{L}_{RNA} addresses the problem of different norms in different domains by re-balancing the average and per-class norms of features in each modality across domains, so that the network can focus on features that are more transferable between domains (Xu et al., 2019). To this end, we include the following term in the RNA formulation:

$$\mathcal{L}_{RNA}^{mod}(m_s, m_t) = \mathcal{L}_{RNA}^g(m_s, m_t) + \mathcal{L}_{RNA}^c(m_s, m_t)$$

where $m \in \{u, v\}$. Combining the three components we have previously defined, the extended RNA formulation in DG settings becomes:

$$\mathcal{L}_{RNA} = \mathcal{L}_{RNA}^g(u_s, v_s) + \mathcal{L}_{RNA}^c(u_s, v_s) \quad (2)$$

and the one for UDA setting is:

$$\begin{aligned} \mathcal{L}_{RNA} = & \mathcal{L}_{RNA}^g(u_s, v_s) + \mathcal{L}_{RNA}^g(u_t, v_t) + \\ & \mathcal{L}_{RNA}^c(u_s, v_s) + \mathcal{L}_{RNA}^c(u_t, v_t) + \\ & \mathcal{L}_{RNA}^{mod}(u_s, u_t) + \mathcal{L}_{RNA}^{mod}(v_s, v_t) \end{aligned} \quad (3)$$

The individual contribution of the three losses is exemplified in Fig. 3. \mathcal{L}_{RNA}^g globally aligns the norms of modalities for each domain. \mathcal{L}_{RNA}^c aligns the norms of modalities per class for each domain. \mathcal{L}_{RNA}^{mod} aligns the norms between domains, separately for each modality. Taken together, the three losses act synergistically. In DG, \mathcal{L}_{RNA}^c supports the work of \mathcal{L}_{RNA}^g , which in turn facilitates the alignment of norms per class to a common value. The addition of \mathcal{L}_{RNA}^{mod} in UDA helps the other two components to ensure that the average and per-class norms of the different modalities are also aligned between source and target.

3.5 Extension to Multiple Modalities

The RNA objective in Eqs. 2 and 3 can be trivially extended to more than two modalities. In DG, the loss can be rewritten as:

$$\mathcal{L}_{RNA} = \mathcal{L}_{RNA}(\mathcal{S}) = \sum_{i=1}^M \sum_{j=i+1}^M \mathcal{L}_{RNA}(i_s, j_s) \quad (4)$$

where i and j span the M modalities. Similarly, the UDA loss becomes:

$$\mathcal{L}_{RNA} = \mathcal{L}_{RNA}(\mathcal{S}) + \mathcal{L}_{RNA}(\mathcal{T}) + \sum_{i=1}^M \mathcal{L}_{RNA}^{mod}(i_s, i_t)$$

where $\mathcal{L}_{RNA}(\mathcal{S})$ and $\mathcal{L}_{RNA}(\mathcal{T})$ are the loss in Eq. 4 for the source and target domains, respectively.

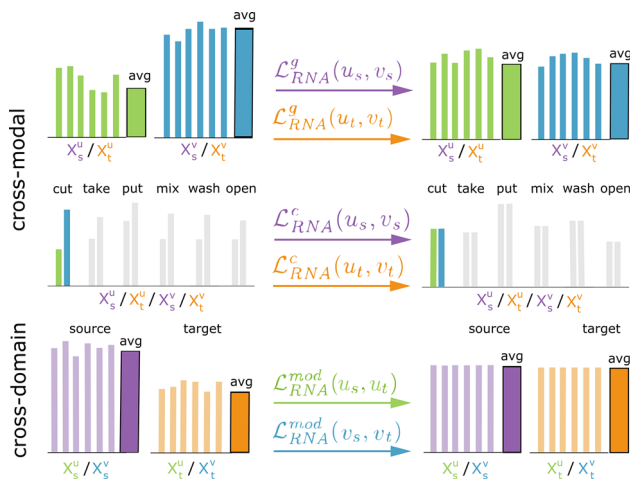


Fig. 3 Individual effects of \mathcal{L}_{RNA} components on feature norms. Each diagram shows norms per class for a single modality and domain (u or v for source or target). **1st row:** \mathcal{L}_{RNA}^g minimizes overall average norms (larger bars on the right) of u and v modalities. **2nd row:** \mathcal{L}_{RNA}^c achieves balanced norms at class level. **3rd row:** \mathcal{L}_{RNA}^{mod} balances class and average norms of the same modality across domains. The diagrams show the norms before (left) and after (right) applying the corresponding \mathcal{L}_{RNA} component

3.6 Learning Objective in UDA

In addition to the loss defined in Eq. 3, to further improve the domain invariant properties of the features (and thus reduce the divergence between domains), we apply an adversarial domain alignment (Ganin & Lempitsky, 2015; Wang et al., 2019). We follow the recipe used in other recent UDA work (Chen et al., 2019; Munro & Damen, 2020; Wei et al., 2022; Jamal et al., 2018), and introduce a classifier that predicts whether features are from the source or the target. This classifier is directly connected to the feature extractors via a Gradient Reversal Layer (GRL) (Ganin & Lempitsky, 2015). The domain classification loss \mathcal{L}_d is then multiplied by a weight λ_d and added to the total loss.

The loss we have introduced so far (i.e., the combination of \mathcal{L}_{RNA} and \mathcal{L}_d) aims to improve the informative and domain invariant properties of the embeddings of the different modalities. However, these two loss components affect the feature extractors F^m and are not back-propagated through the classifier, which therefore only sees the source data and thus has no way to benefit from the target data. The result is that during training, the classifier focuses only on how best to integrate the multi-modal features to improve accuracy in the source domain, and completely ignores the *classification uncertainty* on target.

One approach commonly used in UDA to improve class discrimination in the target domain is to use a mutual information criterion (Bridle et al., 1991) applied to the target data that not only minimizes the prediction uncertainty, but also promotes a uniform distribution of samples between classes.

This is achieved through an Information Maximization (IM) loss defined as the difference between the average entropy of the outputs and the entropy of the average output:

$$\mathcal{L}_{IM} = -\mathbb{E}_{x \in \mathcal{X}_T} \sum_{c=1}^C p_c(x) \log p_c(x) + \sum_{c=1}^C \bar{p}_c \log \bar{p}_c$$

where C is the total number of classes, p_c is the posterior probability for class c , and \bar{p}_c is the mean output score for the current batch.

When we put all the pieces together, we train the model in the UDA setting to minimize the following loss:

$$\mathcal{L} = \mathcal{L}_C + \mathcal{L}_{RNA} + \lambda_d \mathcal{L}_d + \lambda_{IM} \mathcal{L}_{IM}$$

where \mathcal{L}_{RNA} is from Eq. 3 and λ_{IM} is the IM loss weight.

4 Experiments

In this section, we aim to verify the effectiveness of our proposed approach through an empirical evaluation on different multi-modal benchmarks corresponding to a variety of datasets and tasks. These range from action classification (on EPIC-Kitchens-55 Damen et al., 2018, EPIC-Kitchens-100 Damen et al., 2022, and UCF-HMDB Chen et al., 2019) to object recognition (on ROD Lai et al., 2011) and fatigue classification (on CogBeacon Papakostas et al., 2019).

In the analysis, the results are obtained and presented as follows. When a dataset includes different domains, we optimized the models using the average accuracy over all the domain splits reported in the respective experimental protocol. Results were obtained using the same set of hyperparameters for all splits. Therefore, in the following, we excluded from evaluation the methods for which it was obvious (either from the description or from the available source code) that the hyperparameters were optimized for each split.

The rest of the section is organized as follows. We begin by introducing the experimental benchmarks used in our work (Sect. 4.1). The ablation study is given in Sect. 4.2, and the results are presented from Sects. 4.3–4.7.

4.1 Datasets

4.1.1 EPIC-Kitchens-55 (EK55)

This is a large-scale egocentric video benchmark recorded by 32 participants in their own kitchens while performing unscripted activities (Damen et al., 2018). RGB, Audio and Flow data are available in the dataset. To validate our approach, we use the experimental protocol defined in Munro and Damen (2020). According to this protocol, (i) we only

use the three kitchens with the largest amount of annotated samples (hereafter referred to as D1, D2, and D3) and (ii) we consider only verb classification task and a subset of eight labels. The challenges lie not only in the large domain shift that exists between the different kitchens, but also in the unbalanced distribution of classes within and between domains.

4.1.2 EPIC-Kitchens-100 (EK100)

EK100 (Damen et al., 2022) extends EK55 to 45 kitchens, with almost 100h of video and 89,977 annotated action segments. The UDA setting of EK100 is defined as two domains, *Source* (containing labeled training data from 16 participants collected in 2018) and *Target* (i.e., unlabeled videos from the same 16 participants in the same kitchens but collected 2 years later). The segments are annotated with 97 nouns and 300 verbs corresponding to 3369 unique action classes, largely unbalanced and characterised by a long-tailed distribution.

4.1.3 UCF-HMDB

The UCF-HMDB dataset (Chen et al., 2019) was published to study video domain adaptation in third-person action classification. The dataset consists of 3209 videos from the original UCF101 (Soomro et al., 2012) and HMDB51 (Kuehne et al., 2011) datasets, which define the source and target domains used in DG e UDA. The videos are annotated with 12 classes.

4.1.4 ROD

ROD (Lai et al., 2011) is an image-based dataset developed for object recognition tasks. ROD consists of 41,877 samples of 300 everyday objects grouped into 51 categories and captured by an RGB-D camera. ROD is coupled with SynROD (Loghmani et al., 2020), which contains photorealistic renderings from 3D models of the same categories as ROD, and N-ROD (Cannici et al., 2021), which extends both datasets to event modality by introducing real event recordings obtained from ROD samples, as well as simulated events extracted from SynROD's synthetic images. Following the settings proposed in Loghmani et al. (2020), Cannici et al. (2021), this dataset allows the exploration of domain shift between synthetic (source domain) and real data (target domain) in a multi-modal object classification task using RGB-Depth and RGB-Event.

4.1.5 CogBeacon

CogBeacon is a multi-modal dataset collected to analyze the effects of cognitive fatigue on human performance (Papakostas et al., 2019). Volunteers completed three

different computerized versions (V1, V2, and V3) of the Wisconsin Card Sorting Test, a test widely used in experimental and clinical psychology (Lange et al., 2018). Experimental sessions are divided into rounds in which subjects can signal their cognitive fatigue (i.e., sample classes are “fatigue” and “no-fatigue”, with a strong class imbalance towards *fatigue* in split V3). Two modalities are available: (i) EEG data, and (ii) user's movements and facial expressions (recorded by capturing 68 facial keypoints and the face bounding box).

4.2 Ablation Studies

In this section, we present the ablation studies of our approach, all of which have been performed using EK100, as this is the largest and most diverse of all the benchmarks used in our work, thus increasing the statistical significance of these studies.

4.2.1 Experimental Settings

For the sake of clarity, we introduce here the details of the experimental settings used to obtain the results discussed in this Section and in Sect. 4.3.

Evaluation Protocol. We follow the experimental setup for UDA proposed in Damen et al. (2022), where the fine-grained nature of the dataset annotations combined with the large domain and temporal shifts between the source and target domains make the adaptation task very challenging. All the experiments in this section (and in Sect. 4.3) use all three modalities (RGB, Audio, and Flow) available in the dataset. The setting includes a validation split, for which labels are available, and a non-annotated test split. The results of this work are reported on the former, although previous work has also demonstrated the effectiveness of RNA on test data as well (Plizzari et al., 2021; Planamente et al., 2022). Performance is evaluated in terms of Top-1 and Top-5 accuracy of verb and noun predictions and on the combination of the two predictions (action).

Input. RGB, Flow and Audio are processed following Kazakos et al. (2019) by uniformly sampling 25 frames and 1.28 s audio segments along the action. During both training and inference, five of these segments are selected for each modality and fed to the network.

Implementation Details. Frame-level features $f_m \in \mathbb{R}^{25 \times 1024}$ from each modality m are extracted using a TBN architecture (Kazakos et al., 2019) pre-trained on Kinetics (Zisserman et al. 2017) and fine-tuned on the source domain, following the recipe from Damen et al. (2022). Our model is trained on pre-extracted features using this backbone. Five frame features for each segment are uniformly selected and fed to a linear layer, followed by a ReLU activation and dropout with probability 0.5. Frame features are temporally aggregated using a TRN (Zhou et al., 2018) mod-

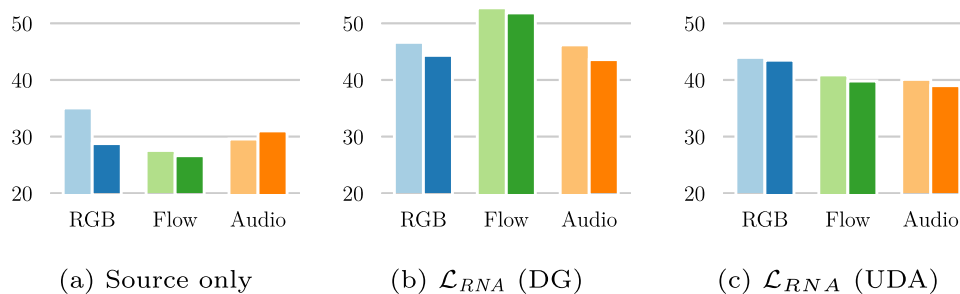


Fig. 4 Verb feature norms across different modalities and settings (DG and UDA). Light (■) and dark colors (■) denote source and target validation domains, respectively. **a** In the “Source Only” setting, different modalities and domains result in unbalanced feature norms. **b** \mathcal{L}_{RNA} in DG improves the alignment between different

modalities, but leaves a gap between the source and target domains. **c** Finally, the contribution of \mathcal{L}^{mod} in \mathcal{L}_{RNA} reduces this gap in UDA, resulting in more consistent feature norms across different modalities and domains

ule to obtain action-level features $f'_m \in \mathbb{R}^{1024}$.¹ To account for the multi-task nature of this setting, we map the features into two components $f'_{m,v}, f'_{m,n} \in \mathbb{R}^{256}$ using a single linear layer, which we call *verb* and *noun features*. These are fed to two separate classifiers to obtain the modality logits for the verb ($y_{m,v}$) and the noun ($y_{m,n}$). Since this benchmark includes a single source and a single target domain, the network is trained for action recognition by applying cross-entropy loss to the sum of *per-modality* logits. We extend RNA to work in this multi-task context by applying the alignment losses separately to the verb and noun features, immediately before the final classifier. Applying the RNA losses to these features ensures that the alignment effect provided by RNA is as close as possible to the classifier, which is heavily influenced by the feature norm values. The network is trained for 30 epochs using a batch size of 128 samples and SGD optimizer with momentum 0.9 and weight decay 10^{-4} . The learning rate is initially set to 0.003 and decreased by a factor of 10 after epochs 10 and 20.

4.2.2 Effects of \mathcal{L}_{RNA} on Norm Alignment

We begin by discussing the contribution of the components of the proposed \mathcal{L}_{RNA} loss. Its goal is to mitigate domain shift issues by balancing the mean feature norms of the different modalities globally (\mathcal{L}_{RNA}^g), at the class level (\mathcal{L}_{RNA}^c), and across domains (\mathcal{L}_{RNA}^{mod}). In the following, we present the results of experiments in which these components are introduced incrementally.

Global alignment: a qualitative analysis. In Fig. 4 we report the mean feature norms for each modality. For simplicity, we will base our discussion on the verb feature norms, since the same observations apply to nouns. In particular, in Fig. 4 we show how the average norms of verb features for

different modalities change on DG and UDA with the contribution of \mathcal{L}_{RNA} .

A preliminary qualitative analysis of the data presented in Fig. 4 shows that \mathcal{L}_{RNA} in DG (Fig. 4b) leads to a better alignment of the average feature norms of the different modalities and to an overall increase of their values with respect to the “Source Only” (Fig. 4a). Recall that the norm formulation in Eq. 2 attempts to solve the alignment task at the batch level, and thus does not guarantee an exact alignment of all average norms. In Fig. 4b, we can also observe the increase in Flow norm in DG compared to “Source Only” (Fig. 4a). Previous studies have shown that Flow is the modality least affected by domain shift in egocentric action recognition (Munro & Damen, 2020), potentially allowing for greater generalization. This could explain why, in DG, the network pays more attention to this modality.

In addition, the availability of target data in UDA enables \mathcal{L}_{RNA} to improve the balance between the norms of the different modalities, so that the model can better use the contributions of each modality to make its final decisions. This improved mutual contribution between modalities (reflected in the increased accuracy reported in Table 2) may explain the (relatively) lower norm of Flow in UDA, which is balanced by increased norms of (i.e., attention to) the other two modalities (RGB and Audio).

Global alignment: a quantitative analysis. To facilitate the assessment of the balancing effect of \mathcal{L}_{RNA} between “Source Only”, DG and UDA norms, we also introduce a quantitative metric. We use the *coefficient of variation* (CV) as a measure of the norm imbalance, with lower CVs indicating more balanced sets of values. CV is defined as follows:

$$CV = \frac{\sigma}{\mu}$$

where σ is the standard deviation and μ is the mean of the observed norm values. The CV values obtained are summarized in Table 1, where, for better clarity, we also report the

¹ Up to this point, the implementation closely follows the official code provided for the EK100 UDA challenge (Damen et al., 2022).

Table 1 Coefficient of variation for DG and UDA feature norms

	CV_S	CV_T	CV_{S+T}
Source Only	0.126	0.076	0.101
\mathcal{L}_{RNA}^g	0.089 (+29.3%)	0.121 (-59.8%)	0.103 (-2.1%)
$\mathcal{L}_{RNA}^g + \mathcal{L}_{RNA}^c$ (DG)	0.075 (+40.4%)	0.098 (-28.7%)	0.081 (+20.1%)
\mathcal{L}_{RNA} (UDA)	0.049 (+61.0%)	0.059 (+22.7%)	0.049 (+50.7%)

CV_S , CV_T and CV_{S+T} are the CVs of the source, target and combined domain(s) respectively. For clarity, we also report the percentages of improvement with respect to the “Source Only” experiment

percentage of improvement (%) with respect to the CV values of the “Source Only”.

As for the average feature norms in DG (Fig. 4b), we have a 40.4% decrease in CV compared to the “Source Only”. It is interesting to note that the application of \mathcal{L}_{RNA}^g alone only contributes to a 29.3% reduction of CV, highlighting the (positive) combined effect of \mathcal{L}_{RNA}^g and \mathcal{L}_{RNA}^c . For the target domain in DG, we can observe that the imbalance between modalities increases (instead of decreasing) by 28.7%, which highlights the need for an alignment loss that works not only between modalities but also between domains.

In UDA, the ability to use the target data contributes to a larger reduction in CV over the “Source Only” on both source (by 61.0%) and target domains (22.7%). When we consider the total imbalance (i.e., we calculate CV considering all source and target values together), CV shows an improvement of 20.1% in DG and of 50.7% in UDA. These values are reflected in progressively greater accuracy in the DG and UDA settings compared to the “Source Only” settings (Table 2).

Class alignment. For assessing the contribution of \mathcal{L}_{RNA}^c , we show in Fig. 5 the evolution of the verb norms of the ten most frequent and the least frequent classes in the DG

settings. In the “Source Only” (Fig. 5a) the per-class mean features norms are largely unbalanced. While the exclusive use of \mathcal{L}_{RNA}^g contributes to a better global balance of the modality norms, it has a small effect on the balancing of the norms per-class (Fig. 5b). On the contrary, when \mathcal{L}_{RNA}^c is also minimized, we can observe a significant improvement of their alignment (Fig. 5c).

These qualitative observations are also reflected in the CV metric computed on the class norms. Indeed, the use of \mathcal{L}_{RNA}^g leads to a minor improvement in “Source Only” CV (37.8% and 19.5%, respectively, for source and target features) compared to that obtained by the combination of \mathcal{L}_{RNA}^g and \mathcal{L}_{RNA}^c (62.5% and 49.1%).

Overall effect on feature norms. To give further insight into the impact of \mathcal{L}_{RNA} , we show in Fig. 6 a scatter plot of the validation set in DG. This diagram is obtained by plotting the RGB, Flow and Audio feature norms of each sample in a 3D space whose axes are the norms of the three modalities. To make the plot easier to read, rather than using a single 3D representation, we present it as three separate sections along the coordinate planes defined by the feature pairs. The goal of these visualizations is to illustrate the changes in the shape of the resulting manifold.

It can be seen that the “Source Only” features are widely distributed and correspond to a manifold with a largely irregular shape. This is due to misalignment between the feature norms of the different modalities. When the \mathcal{L}_{RNA} loss is applied, the manifold becomes more spherical and compact, reflecting the improvement in the alignment of the modality norms. It is also possible to note an increase in the average feature norm values that moves the manifold towards the upper right region of the 2D dimensional plots.

Table 2 Ablation on different loss components

	Verb@1	Noun@1	Action@1	Δ Acc.
Source only	46.79	26.79	18.29	–
<i>DG</i>				
\mathcal{L}_{RNA}^g	49.53	27.50	18.91	1.36
\mathcal{L}_{RNA}^c	50.51	27.75	19.44	1.94
$\mathcal{L}_{RNA}^g + \mathcal{L}_{RNA}^c$	50.75	27.92	19.81	2.20
<i>UDA</i>				
\mathcal{L}_{RNA}^g	49.98	27.79	19.44	1.78
$\mathcal{L}_{RNA}^g + \mathcal{L}_{RNA}^c$	50.46	28.49	19.77	2.28
$\mathcal{L}_{RNA}^g + \mathcal{L}_{RNA}^c + \mathcal{L}_{RNA}^{mod}$	49.94	29.48	19.87	2.48
$\mathcal{L}_{RNA} + \mathcal{L}_d$	50.59	<u>29.38</u>	<u>20.04</u>	2.71
$\mathcal{L}_{RNA} + \mathcal{L}_{IM}$	51.04	28.86	19.97	2.67
$\mathcal{L}_{RNA} + \mathcal{L}_d + \mathcal{L}_{IM}$	<u>50.82</u>	29.19	20.05	2.73

Δ Acc. is the average accuracy improvement for the verb, noun and action metrics. Best in **bold** and the runner-up underlined

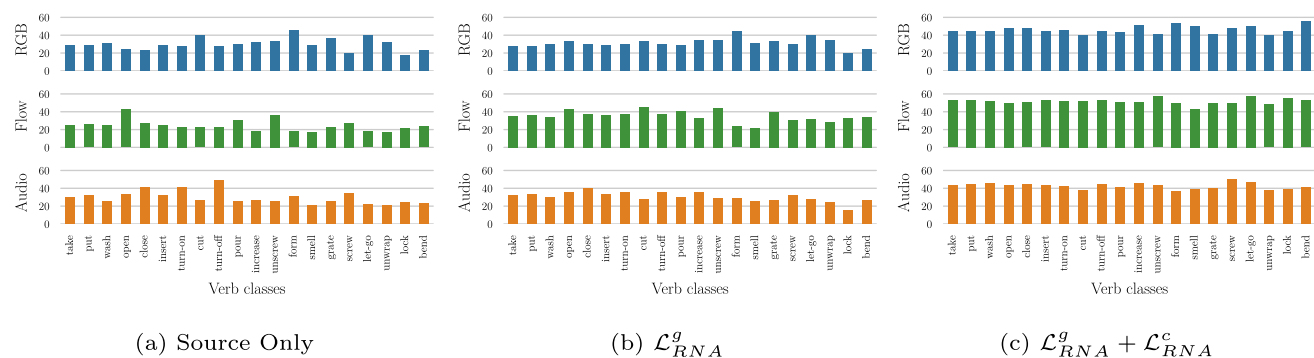


Fig. 5 Feature norms of the top 10 most and least common classes from the target validation split of EPIC-Kitchens-100. While \mathcal{L}_{RNA}^g improves the alignment of different modalities, there is still an imbalance between

classes. The addition of the per-class variant of RNA greatly improves this alignment, resulting in more uniform feature norms across different classes

4.2.3 Effect of Loss Components

Table 2 details the contribution of the different loss components to the final performance in both DG and UDA settings. For better evaluation, we also show the average improvement in Top-1 accuracy of verb, noun, and action with respect to “Source Only” (Δ Acc.). The combination of global and class components in DG ($\mathcal{L}_{RNA}^g + \mathcal{L}_{RNA}^c$, Δ Acc. = 2.20) improves accuracy over \mathcal{L}_{RNA}^g and \mathcal{L}_{RNA}^c alone (1.36 and 1.94, respectively), showing that the combination of the two components effectively reduces domain shift. Indeed, while \mathcal{L}_{RNA}^g aligns modalities globally, possibly penalizing minority classes in unbalanced distributions, \mathcal{L}_{RNA}^c enforces alignment for each class individually. The ability to use target data in UDA boosts the accuracy improvement to 1.78 for \mathcal{L}_{RNA}^g and 2.28 for $\mathcal{L}_{RNA}^g + \mathcal{L}_{RNA}^c$, with \mathcal{L}_{RNA}^{mod} further contributing to reach an average improvement of 2.48.

As explained in Sect. 3.6, the learning objective in the UDA setting also benefits from two other losses, namely the adversarial domain loss \mathcal{L}_d , which aims to improve the transferability of features across domains, and the Information Maximization loss \mathcal{L}_{IM} , which aims to minimize the classification uncertainty between target classes. \mathcal{L}_d provides a stronger improvement in this particular case (2.71), while \mathcal{L}_{IM} has a minimal effect on the overall accuracy. However, we note that the mutual contribution of the latter two terms (\mathcal{L}_d and \mathcal{L}_{IM}) also depends on the task and benchmark considered, as other experiments show more pronounced benefits for \mathcal{L}_{IM} .

4.2.4 Multi-modal Adaptation Capabilities

Another interesting question is whether the proposed method allows effective integration of multiple modalities in the final decision and whether the use of multiple modalities also helps to improve the domain adaptation capabilities of the model.

Table 3 summarizes the results obtained comparing experiments with modality pairs and with all three modalities. It shows that the latter not only outperforms all other modality pairs in terms of results, but also shows better generalization properties, showing an improved delta compared to its “Source only” (2.73) compared to 2.06, the best two-modality improvement obtained with Flow + Audio. These results suggest that our method is effective in combining the different modalities to improve the overall accuracy and the generalizability of the features obtained.

4.2.5 Modality Drop

In Table 4, we present an experiment to investigate the impact of modality imbalance during training. In particular, we investigate the scenario in which a modality is “unexpectedly” lost at inference time, without a training strategy accounting for this possibility. This scenario, also presented in Gong et al. (2023), is relevant because there may be constraints at inference time, such as power, computational or privacy constraints, or an anomaly of an input device that prevent the use of all modalities.

The basic idea of our approach is to help the model learn equally from the different modalities by integrating their contribution. While it is clear that the unexpected loss leads to a drop in accuracy, we can also expect that the effect of RNA is to make the model more robust to such a modality drop than the “Source Only” model, since the latter is less able to exploit the synergies between modalities and, thus, more vulnerable to dominant modalities. This expectation is confirmed by the results in Table 4, which are consistent with the observations of Gong et al. (2023), and show different but consistent effects on “Source Only” when different modalities are dropped at test time (i.e., large accuracy drops compared to the results in Table 2). At the same time, these results show that the balancing effect of RNA can potentially

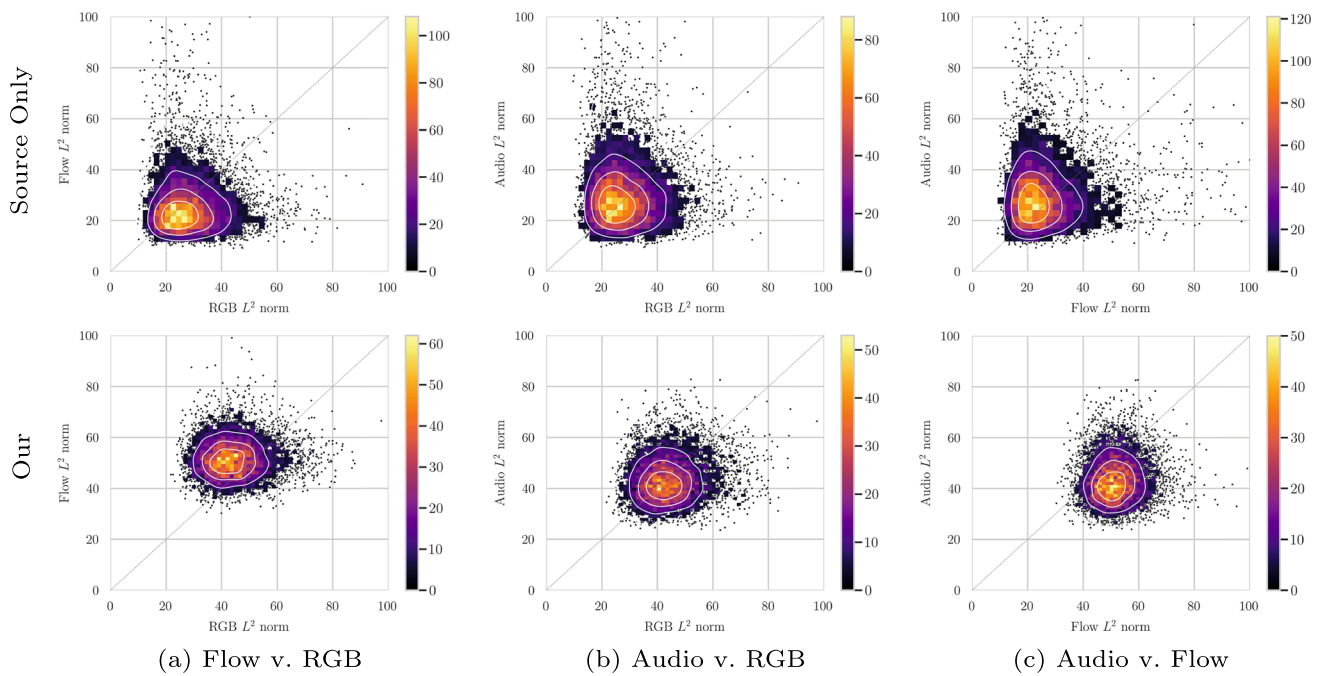


Fig. 6 Comparison of the feature norms before (top) and after (bottom) application of \mathcal{L}_{RNA}^g and \mathcal{L}_{RNA}^c . The dots represent the samples in the validation dataset. The color bar on the right represents increasing density values. The original features, i.e. “Source Only”, show a wide range

of values and an irregular shape, reflecting the misalignment between the features norms of the two modalities. The RNA loss re-balances the two, as evidenced by the more globular distribution while also shifting the average norms towards higher values

Table 3 Top-1 classification accuracies (%) on modality pairs on EPIC-Kitchens-100 (Damen et al., 2022)

	Verb@1	Noun@1	Action@1	Δ Acc.
<i>RGB + Flow</i>				
Source only	44.80	25.35	16.33	–
Our (DG)	45.95	26.65	16.94	1.02
Our (UDA)	47.64	26.49	16.91	1.52
<i>RGB + Audio</i>				
Source only	39.91	24.18	14.84	–
Our (DG)	42.04	25.54	15.67	1.44
Our (UDA)	42.26	26.45	15.98	1.92
<i>Flow + Audio</i>				
Source only	45.11	21.98	15.37	–
Our (DG)	48.87	23.44	16.49	2.12
Our (UDA)	48.42	23.51	16.71	2.06
<i>RGB + Flow + Audio</i>				
Source only	46.79	26.79	18.29	–
Our (DG)	50.75	27.92	19.81	2.20
Our (UDA)	50.82	29.19	20.05	2.73

Δ Acc. is the average accuracy improvement for the verb, noun and action metrics

help the model reduce the impact of the lost modality, as it can take advantage of a better mutual contribution from the remaining ones.

Table 4 Modality drop

	Verb@1	Noun@1	Action@1	Δ Acc.
<i>No audio @ test</i>				
Source only	41.61	21.91	13.07	–
DG	44.03	24.44	14.89	2.26
UDA (\mathcal{L}_{RNA})	44.08	24.77	15.25	2.50
<i>No flow @ test</i>				
Source only	30.58	20.33	10.63	–
DG	36.88	22.82	12.89	3.69
UDA (\mathcal{L}_{RNA})	36.67	21.83	12.46	3.14
<i>No RGB @ test</i>				
Source only	37.69	17.99	12.41	–
DG	46.70	18.92	13.53	3.69
UDA (\mathcal{L}_{RNA})	46.51	19.37	13.55	3.78

All configurations are trained on all input modalities. At inference time, we simulate the loss of a modality, resulting in large performance drops that RNA helps mitigate

4.3 Experiments on EK100

Unlike the following benchmarks, where we describe the experimental protocol, inputs, implementation details, and baselines used to evaluate the results, in this section we present only the baselines for EK100, as the previous elements were introduced in Sect. 4.2.1.

Table 5 Classification accuracies (%) on EPIC-Kitchens-100 (Damen et al., 2022)

Methods	Network	Verb@1	Noun@1	Action@1	Verb@5	Noun@5	Action@5	Δ Acc.
<i>DG</i>								
Source only	TBN-TRN	47.14	27.35	18.99	75.27	49.36	41.82	–
MM-SADA (SS) (Munro & Damen, 2020)	TBN-TRN	47.76	27.93	19.15	77.07	49.77	42.90	0.45
Source only	TBN-TRN	<u>50.27</u>	29.04	19.96	<u>81.74</u>	52.14	46.74	–
GB (Wang et al., 2020)	TBN-TRN	50.18	29.60	<u>20.26</u>	81.82	<u>52.57</u>	46.86	0.26
Source only	TBN-TRN	46.79	26.79	18.29	75.39	48.44	41.36	–
Our (DG)	TBN-TRN	50.75	27.92	19.81	80.64	51.37	45.33	2.20
Source only [†]	TBN-TRN	49.81	28.55	19.77	81.10	51.90	46.22	–
Our [†] (DG)	TBN-TRN	50.20	<u>29.31</u>	20.30	81.58	52.68	<u>46.76</u>	0.56
<i>UDA</i>								
Source only	TBN-TRN	46.70	27.78	19.20	75.42	48.27	42.12	–
TA ³ N (Chen et al., 2019)	TBN-TRN	<u>48.44</u>	28.87	19.61	75.95	50.12	43.36	1.08
Source only	TBN-TRN	47.14	27.35	18.99	75.27	49.36	41.82	–
MM-SADA (Munro & Damen, 2020)	TBN-TRN	<u>48.44</u>	28.26	19.25	<u>77.56</u>	<u>50.59</u>	<u>43.41</u>	0.82
Source only	TBN-TRN	47.69	28.48	19.61	–	–	–	–
CIA (Yang et al., 2022)	TBN-TRN	48.34	29.50	20.30	–	–	–	0.79
Source only	TBN-TRN	46.79	26.79	18.29	75.39	48.44	41.36	–
Our (UDA)	TBN-TRN	50.82	<u>29.19</u>	<u>20.05</u>	80.89	52.18	46.04	2.73

Results are reported in terms of Top-1 and Top-5 classification accuracy for the noun, verb and action metrics. Δ Acc. is the average Top-1 accuracy improvement. [†]These experiments are trained using the cross entropy loss on both the fused logits as well as on the *per-modality* logits. Best in **bold**, runner-up underlined

4.3.1 Baselines

We compare our method with MM-SADA (Munro & Damen, 2020), TA³N (Chen et al., 2019), and CIA (Yang et al., 2022). As for MM-SADA, the original approach works only with RGB and Flow modalities. Therefore, to integrate the Audio modality, we use two separate branches, one for RGB-Flow and the other for RGB-Audio modalities (as in Planamente et al. (2022)). The adversarial branch is applied individually to each modality.

4.3.2 Results

Results are given as Top-1 and Top-5 accuracy for verb, noun, and action (Table 5). For each baseline, we also report the relative “Source Only”, average improvement in terms of Top-1 accuracy. For the DG setting, we compare our approach to two alternative methods. The first is MM-SADA (SS), a modified version of MM-SADA, which applies only the original self-supervised alignment task to the source domain modalities and does not consider the adversarial alignment component of the method (which requires target data). The second approach is Gradient Blending (GB), which attempts to find an optimal mixture of modalities according to their overfitting behavior. Such a mixture is achieved by combin-

ing a cross-entropy loss for each modality and a loss for their fusion with appropriate weights.²

Analyzing the accuracy across different labels, we observe that GB performs best, while our approach ranks as the runner-up and MM-SADA (SS) lags slightly behind. However, when considering the improvements relative to the “Source Only” baseline, our method shows higher delta accuracy compared to GB. This result seems to indicate that our method makes a more significant contribution to reducing the domain shift. We also find the approach proposed in Wang et al. (2020) interesting as it shares similarities with our method in terms of improving the balance between modalities for better classification accuracy. To investigate this further, we perform additional experiments by applying our method to the “Source Only” results obtained from Gradient Blending, i.e., using multiple classification losses but without reweighting them. These additional experiments are indicated with a [†] symbol. The results shown in Table 5 are promising. Our method achieves the best action accuracy and is competitive with GB (and also comparable with CIA, the state-of-the-art in UDA). It is important to note that our standard solution

² The original version of GB uses only RGB and Audio. The optimal weights for combining losses were taken from Damen et al. (2020), and the weight for the missing component, i.e., Flow, was tuned appropriately for this work.

addresses the alignment problem with an adaptive approach that, unlike GB, is independent of the model and dataset used and requires only two hyperparameters: λ_g and λ_c .

In the UDA experiments, we observe that although our method ranks second in terms of action accuracy, it has better delta accuracy improvements compared to all other competitors. Furthermore, the results on other evaluation metrics are comparable to those of the other proposed baselines. It is noteworthy that a significant portion of the improvements can be observed in the DG phase, where the target domain is not accessed. This observation highlights the generalization advantage of RNA in coping with domain shifts.

4.4 Experiments on EK55

4.4.1 Evaluation Protocol

We adopt the experimental protocol of Munro and Damen (2020) and evaluate performance in a single-source setting ($D_i \rightarrow D_j$) on the three domains described in Sect. 4.1. Despite the small size of this setting compared to EK100, it remains a highly valued and challenging benchmark in the field of egocentric action recognition due to the large domain shift between these domains and the unbalanced label distribution. In the experiments, we restrict our analysis to the RGB+Flow and RGB+Audio modality combinations, which are the ones recent work in the literature focus on.

4.4.2 Baselines

We compare our results with several state-of-the-art UDA methods. The first group (GRL Ganin et al., 2016, MMD Long et al., 2015, AdaBN Li et al., 2018, and MCD Saito et al., 2018) includes approaches originally developed as image-based methods and later adapted to work with video inputs. The second group includes more recent methods such as MM-SADA (Munro & Damen, 2020), the contrastive-based methods proposed in Kim et al. (2021) and Song et al. (2021) (STCDA), and the recently published CIA (Yang et al., 2022). In our comparison, we use the results reported in the original paper for each baseline.

4.4.3 Input

As for the input, different sampling strategies are used to allow a fair comparison with the existing baselines. When using *dense sampling*, a clip of 16 consecutive frames is randomly sampled from the video. When using *uniform sampling*, 16 frames evenly distributed over the video are sampled. At test time, the same sampling strategy is used as in training, except that five clips are fed into the network instead of one, as suggested in Wang et al. (2016), and the predictions are averaged. Following the experimental setting from Munro

and Damen (2020), during training, random clipping, scale shifts, and horizontal flipping are used for data augmentation, while in testing, only central cropping is applied. As for the aural information, we follow Kazakos et al. (2019) and convert the audio track into a 256×256 matrix representing the log spectrogram of the signal. The audio clip is first extracted from the video and sampled at 24 kHz. Then, the Short-Time Fourier Transform (STFT) is calculated with a window length of 10 ms, a skip size of 5 ms, and 256 frequency bands. For the Flow input, we use the same sampling strategy as for RGB.

4.4.4 Implementation Details

Both the RGB and Flow streams use an I3D model (Carreira & Zisserman, 2017) pretrained on Kinetics (Zisserman et al., 2017), following the experimental setting from Munro and Damen (2020). Following Kazakos et al. (2019), the audio feature extractor uses the BN-Inception model (Ioffe & Szegedy, 2015) pre-trained on ImageNet (Deng et al., 2009). The feature extraction backbones are trained end-to-end. For each modality m , features have shape $f_m \in \mathbb{R}^{1024}$. Logits are computed separately for each modality using a linear layer and summed. We train the network for 5000 iterations using the SGD optimizer using momentum 0.9 and weight decay 10^{-7} . The learning rate for RGB and Flow is set to 0.001 and reduced to 2×10^{-4} at step 3000, while for Audio the learning rate is set to 0.001 and decremented by a factor of 10 at steps {1000, 2000, 3000}. The batch size is set to 128.

4.4.5 Results

We begin by discussing the UDA results, which are summarized in Table 6. Given the relevance of sampling strategies in the video context, especially for the RGB+Flow combination (Chen et al., 2021), we divide the results into different sections based on the sampling used for each modality: dense (D) or uniform (U). Most baselines use dense sampling (D–D), while CIA is the only method that uses uniform sampling (U–U) for both modalities. In both cases, we compare the baselines to our UDA method using the same sampling strategy.

The results show that CIA with uniform sampling outperforms the dense sampling-based methods. This observation confirms the findings in Chen et al. (2021), which emphasizes that uniform sampling usually allows the network to learn more information. We also observe that our UDA approach achieves state-of-the-art results for both dense and uniform samplings.

To further confirm the importance of sampling, we conduct experiments with a mixed sampling strategy (i.e., D for RGB and U for Flow, Table 6). Since none of the baselines use this sampling, we only present our results for the

Table 6 Classification accuracies (%) on EPIC-Kitchens-55 (Damen et al., 2018), using the evaluation protocol from Munro and Damen (2020), divided by modalities

Method	Sampling	D1→D2	D1→D3	D2→D1	D2→D3	D3→D1	D3→D2	Mean
<i>RGB + Flow</i>								
Source only	D–D	42.00	41.20	42.50	46.50	44.30	56.30	45.47
GRL (Ganin et al., 2016)	D–D	50.20	44.70	46.90	50.80	50.20	53.60	49.40
MMD (Long et al., 2015)	D–D	46.60	39.20	43.10	48.50	48.30	55.20	46.82
AdaBN (Li et al., 2018)	D–D	47.00	40.30	44.60	48.80	47.80	54.70	47.20
MCD (Saito et al., 2018)	D–D	46.50	43.50	42.10	51.00	47.90	52.70	47.28
DAAA (Jamal et al., 2018)	D–D	50.00	43.50	46.50	51.50	51.00	53.70	49.37
MM-SADA (Munro & Damen, 2020)	D–D	49.50	44.10	48.20	52.70	50.90	56.10	50.25
Kim et al. (2021)	D–D	50.30	46.30	49.50	52.00	51.50	56.30	50.98
STCDA (Song et al., 2021)	D–D	52.00	45.50	49.00	52.50	52.60	55.60	<u>51.20</u>
Our (UDA)	D–D	50.84	47.14	48.86	54.38	50.6	58.43	51.71
Source only	U–U	43.20	42.50	43.0	48.0	43.0	55.50	45.90
CIA (Yang et al., 2022)	U–U	52.50	47.80	49.80	53.20	52.20	57.60	<u>52.18</u>
Our (UDA)	U–U	52.84	47.49	54.41	54.11	55.53	61.64	54.34
Source only	D–U	54.25	50.72	54.87	56.41	51.65	61.27	54.86
Our (DG)	D–U	56.00	50.39	56.25	56.37	56.73	61.63	<u>56.23</u>
Our (UDA)	D–U	57.33	52.84	57.19	56.78	57.27	62.03	57.24
<i>RGB + Audio</i>								
Source only	D–D	39.03	39.17	35.27	47.52	40.255	49.98	41.87
GRL (Ganin et al., 2016)	D–D	41.02	43.04	39.36	49.25	38.77	50.56	43.67
MMD (Long et al., 2015)	D–D	42.40	43.84	40.87	48.13	41.46	50.03	44.46
AdaBN (Li et al., 2018)	D–D	36.64	42.57	33.97	46.63	40.51	51.2	41.92
MM-SADA (Munro & Damen, 2020)	D–D	48.90	46.66	39.51	50.89	45.42	55.14	<u>47.75</u>
Our (DG)	D–D	42.55	41.77	42.73	51.09	42.63	54.24	46.21
Our (UDA)	D–D	46.65	47.22	46.18	52.30	44.04	56.18	48.76

Results are grouped by the sampling strategy used for a fair comparison. Best in **bold**, runner-up underlined

“Source Only”, DG, and UDA methods. We note that the “Source Only” method already achieves remarkable results (up to 3% better than our method with uniform sampling), which are further improved in both DG and UDA (despite the smaller difference with the “Source Only” compared to other samplings). One possible explanation for the improved performance with mixed sampling is that it allows for better exploitation of the distinct properties of the two modalities. Dense sampling facilitates a more accurate characterization of static appearance information (RGB) over a short temporal range, while uniform sampling enables the use of a wider temporal range to capture the dynamic information conveyed by Flow.

When combining RGB and Audio modalities, our UDA approach consistently achieves the best results (7% improvement over “Source Only” and 1% improvement over the state-of-the-art method). This result confirms the potential of our method, even when dealing with the fusion of heterogeneous modalities.

Finally, we discuss the results we obtained in DG for both modality combinations. For RGB+Flow, we report the results obtained with the mixed sampling strategy (D–U), i.e., the sampling that yields the best performance. In the DG setting, our method improves the “Source Only” by up to 2% and 5% for RGB+Flow and RGB+Audio, respectively. Furthermore, the performance obtained in the DG setting is comparable to that of the UDA setting, with a deviation of -1.01% and -2.55% for RGB+Flow and RGB+Audio, respectively. Although no other DG methods are available for comparison in this context, these results show that the DG setting can compete with several existing UDA methods that benefit from target data during training.

4.5 Experiments on UCF-HMDB

4.5.1 Evaluation Protocol

We follow the same experimental setting proposed in Zhou et al. (2018), which includes the $U \rightarrow H$ and $H \rightarrow U$ shifts

in a multi-modal setting that includes the RGB and Flow modalities available with this dataset.

4.5.2 Input

For both RGB and Flow, the training input consists of 16 consecutive frames with resolution 224×224 pixels. In testing, we use five clips uniformly sampled across the video and average the predictions. We use the same data augmentations as described in Sect. 4.4 for EK55.

4.5.3 Baselines

We compare our approach with various multi-modal UDA approaches (MM-SADA Munro & Damen, 2020, STCDA Song et al., 2021, the method of Kim et al., 2021 and CIA Yang et al., 2022). To allow a fair comparison, all multi-modal results are based on the same backbones and the same pre-training.

4.5.4 Implementation Details

The backbone for both RGB and Flow is an I3D pre-trained on Kinetics (Zisserman et al., 2017). The learning rate is set to 0.01 and we train the model for 20 epochs with batch size of 32. We use SGD as the optimizer with a momentum of 0.9 and a weight decay of 10^{-7} .

4.5.5 Results

In Table 7, we present the classification accuracy of our method and several baselines. To ensure a fair comparison, we report the results of the ‘‘Source Only’’ model from the original paper for all baselines.

In absolute terms, our approach achieves very competitive performance under the UDA setting and is the second best of all methods in terms of accuracy, outperforming all baselines except CIA. However, the better ‘‘Source Only’’ result of CIA, which was difficult to reproduce in our experiments, should be emphasized. This result could be attributed to its particular architectural design choices and the integration of spatial consensus between RGB and Flow modalities. However, it should be noted that the method proposed by CIA cannot be easily extended to other modalities, as shown by their work on integrating the Audio modality in EK100 (Yang et al., 2022). In contrast, our approach provides a more versatile and adaptable solution that can be applied to different modalities without significant architectural changes.

Furthermore, our approach shows remarkable domain shift reduction capabilities. When we compare the performance gains of our method with other baselines, we observe that our approach achieves improvements over the ‘‘Source Only’’ baseline that are comparable to those obtained by

Table 7 Classification accuracies (%) on UCF-HMDB on RGB+Flow combination

Method	U→H	H→U	Mean
Source only	82.8	90.7	86.7
MM-SADA (Munro & Damen, 2020)	84.2	91.1	87.6
Source only (Song et al., 2021)	82.8	89.8	86.3
STCDA (Song et al., 2021)	83.1	92.1	87.6
Source only (Kim et al., 2021)	82.8	90.7	86.7
Kim et al. (2021)	84.7	92.8	88.7
Source only (Yang et al., 2022)	86.1	92.5	89.3
CIA (Yang et al., 2022)	88.3	94.1	<u>91.2</u>
Source only (conc) (Yang et al., 2022)	85.8	93.5	89.5
CIA (conc) (Yang et al., 2022)	90.6	94.2	92.4
Source only	83.6	94.1	88.9
Our (DG)	83.3	94.9	89.1
Our (UDA)	86.4	94.3	90.4

Best in **bold**, runner-up underlined

other methods. For example, MM-SADA, STCDA, Kim et al. (2021), and CIA show gains of 0.9%, 1.3%, 2% and 1.9%, respectively, while our approach achieves a gain of 1.5%, with a maximum improvement of up to 3% in the U → H shift. This highlights the effectiveness of our method in adapting to target domains and mitigating the negative effects of domain shift.

4.6 Experiments on ROD

4.6.1 Evaluation Protocol

We follow the experimental protocol in Loghmani et al. (2020) for RGB-depth modalities, and the one in Cannici et al. (2021) for RGB-event. The studied shift is a synthetic-to-real domain shift, with synthetic source data and real target data (SynROD → ROD). RGB and depth modality in the synthetic domain are rendered, while events in the synthetic domain are simulated using ESIM (Rebecq et al., 2018).

4.6.2 Baselines

We compare our results with standard image-based UDA methods, namely GRL (Ganin et al., 2016), MMD (Long et al., 2015), SAFN (Xu et al., 2019) and Entropy (Wu et al., 2021), which we extend to operate on multiple modalities. We also compare with Relative Rotation (Loghmani et al., 2020), a method specifically designed to operate on multiple modalities. It consists in a self-supervised task asking the network to predict the relative rotation between two modalities of the same input sample, e.g., an RGB and a depth image.

Table 8 SynROD→ROD accuracy (%) results

	RGB + D	RGB + E
Source only	47.70	49.19
GRL (Ganin et al., 2016)	59.51	55.11
MMD (Long et al., 2015)	62.57	62.39
SAFN (Xu et al., 2019)	62.40	<u>66.87</u>
Entropy (Wu et al., 2021)	63.12	66.23
Relative rotation (Loghmani et al., 2020)	<u>66.68</u>	66.68
Our (DG)	50.06	50.61
Our (UDA)	82.36	78.52

Best in **bold**, runner-up underlined

4.6.3 Input

Event representations, depth images and RGB images are pre-processed and augmented during training following the procedure in Loghmani et al. (2020). Depth images are colorized with surface normal encoding, as in Aakerberg et al. (2017). Input images are normalized with the same mean and variance used for the ImageNet pre-training, while we kept event representations un-normalized as this provided better performance. We use voxelgrid representation for events with 9 bins as in Cannici et al. (2021).

4.6.4 Implementation Details

All backbones are implemented using ResNet-18 (He et al., 2016), pre-trained on ImageNet (Deng et al., 2009). All the parameters of the network, including the pre-trained parameters, are updated during training, as in Loghmani et al. (2020). We train all network configurations using SGD as optimizer, batch size 64 and weight decay 0.003.

4.6.5 Results

Table 8 provides a comparison of our method with different baselines for the SynROD→ROD adaptation task using RGB, depth, and event modalities.

In UDA setting, our method shows remarkable performance gains over all baselines for both RGB+Depth and RGB+Event combinations. For RGB+Depth, our approach achieves improvements of up to 20% over the baselines. These results demonstrate the capability of our method in reducing the domain shift and improving classification accuracy when adapting from the synthetic domain (SynROD) to the real-world domain (ROD). Our method also achieves improvements of up to 10% over the baseline values for the RGB+Event combination. This is a further evidence of the effectiveness of our method in handling the domain shift and improving classification accuracy even in the presence of event data.

However, for DG, the performance gains are relatively smaller compared to the *Source Only* model for both modality combinations. This can be attributed to the inherent challenges of the synthetic-to-real setting, where there is a significant gap between the feature distributions of the source (SynROD) and target (ROD) domains. The unavailability of target data during training limits the generalization capabilities of the model, resulting in modest improvements over *Source Only*. In contrast, in UDA setting, our method achieves better performance by effectively using the target domain information to bridge the domain gap.

Thus, we can conclude that the success of our method in mitigating the shift from synthetic to real domains highlights its potential for various fields and applications, such as robotics, autonomous driving, and augmented reality, where synthetic training data are largely used.

4.7 Experiments on CogBeacon

4.7.1 Evaluation Protocol

We follow the experimental protocol in the supplemental of Lv et al. (2021), evaluating the performance in the single-source setting ($V_i \rightarrow V_j$) using three different domains (V1, V2, and V3), for a total of six splits.

4.7.2 Baselines

We compare our results with those in Lv et al. (2021) (in particular with DLMM, the Differentiated Learning framework proposed in Lv et al., 2021) and with those obtained in our experiments with different UDA methods, namely SAFN (Xu et al., 2019), GRL (Ganin et al., 2016), MMD (Long et al., 2015), and MM-SADA (Munro & Damen, 2020). These two lists of results are presented separately in Table 9 because the number of samples does not match that used in Lv et al. (2021) (i.e., we have 2,240, 2,432, and 2,300 for domains V1, V2, and V3, respectively).

4.7.3 Input

The EEG signals are characterized using a total of 24 temporal and spectral features (see Papakostas et al., 2019 for details). The face data are represented as a vector combining the average values of the face data and their standard deviation, yielding a total of 280 values.

4.7.4 Implementation Details

Both backbones are implemented by three 1D convolutional blocks with kernel size three and stride one, followed by a MaxPool layer and ReLU as the activation function. The output channels are 16, 32 and 64 for EEG signals and 8, 16 and 32 for the face keypoint model. The latter ends with a fully

Table 9 CogBeacon accuracy (%) results

Results from Lv et al. (2021)		Ours	
Source only	63.64	Source only	61.80
MDANN (Qi et al., 2018)	66.83	SAFN (Xu et al., 2019)	64.01
MCD (Saito et al., 2018)	66.75	GRL (Ganin et al., 2016)	64.24
CBST (Zou et al., 2019)	67.71	MM-SADA (Munro & Damen, 2020)	65.40
MM-SADA (Munro & Damen, 2020)	<u>67.75</u>	MMD (Long et al., 2015)	<u>65.58</u>
DLMM (Lv et al., 2021)	70.47	Our (DG)	62.64
		Our (UDA)	68.63

Best in **bold**, runner-up underlined

connected layer with an output of 64 to match the output of the EEG backbone. Model weights are randomly initialized. Predictions for each modality are computed with a single FC layer followed by a LogSoftmax. We train the model for 90 epochs using Adam as the optimizer. In all experiments, the learning rate was set to $1e - 3$ and decremented by a factor of 10 after 70 epochs.

4.7.5 Results

In Table 9, we present the classification accuracy of our method and several baselines for the CogBeacon dataset.

When we compare our approach to the “Source Only” baseline, we observe a significant improvement in the UDA setting. Our UDA accuracy of 68.63% significantly outperforms the “Source Only’s” accuracy of 61.80%. On the other hand, the improvement in the DG setting is not so significant, with an accuracy of 62.64%. However, even in this setting, we observe a potential for softening the domain shift and achieve better performance.

It is worth noting that, on our settings, our method outperforms the other UDA methods used for comparison. Among the UDA baselines, MMD achieves the highest accuracy of 65.58%, which is 3.05% less than our method. We are also competitive with more complex domain matching approaches such as CBST, which involve the generation of pseudo-labels or the use of confidence-based selection strategies (Zou et al., 2019).

However, compared to the results reported in Lv et al. (2021), our method falls behind DLMM, which achieves 70.47% accuracy, significantly better than our results. Nevertheless, a more detailed analysis reveals interesting insights when considering the improvements over the “Source Only” baseline. DLMM achieves an improvement of 6.83%, while our method has an improvement of 7.03%, which can be considered equivalent. That said, it is worth mentioning that our approach is characterized by its simplicity compared to DLMM. DLMM requires multiple training stages and uses a more complex curriculum learning approach with teacher/student models for different modalities. In contrast, our method is lighter and simpler to train, making it a more practical option for real-world applications.

4.8 Limitations

The proposed approach provides interesting performance in many cases, as shown by our experiments with a variety of tasks and scenarios. While the simplicity of the method is certainly a strength, it may be viewed as less effective when compared to methods developed and tuned for a specific task and benchmark. However, we believe that this limitation does not undermine the overall effectiveness of the proposed approach, as it provides a viable alternative for addressing various tasks without requiring significant computational resources or architectural changes.

Another limitation we observed arises from the fact that in many real-world cases the data distributions are strongly unbalanced, leading to lower precision for the tail classes (Buda et al., 2018). The literature shows how this imbalance translates into unbalanced norms of classification weights per class (Guo & Zhang, 2017; Kim & Kim, 2020) as well as unbalanced norms of features per class (Wu et al., 2017; Li et al., 2022). In developing our method, we expected that balancing the norms per class could have a positive effect also in rebalancing the weights of the classifier for the tail classes. However, our experimental results show that this effect is not present. This opens up possibilities for future developments to incorporate this objective into RNA as an additional component that rebalances the weights of the classifier.

5 Conclusion

This work introduces a novel approach to address the problem of multi-modal domain adaptation. Our method is motivated by the observation that differences in the marginal distributions of modalities can significantly affect the training process, leading to suboptimal performance and imbalances in feature norms. To tackle these issues, we introduced the Relative Norm Alignment (RNA) loss, which aims to balance the norms of features extracted by the network across different domains and modalities to improve overall accuracy. This loss is combined with adversarial domain loss and Information Maximization in UDA settings to

enhance feature transferability and regularization in the target domain. Our experimental results have shown that the proposed RNA approach can either outperform or compete with several state-of-the-art methods in various multi-modal classification tasks, demonstrating its effectiveness and flexibility. Most notably, our approach is characterized by its simplicity and lightweight nature, allowing it to be easily integrated into different architectures and contexts without requiring complex modifications. This inherent adaptability makes RNA a promising solution for real-world applications where multi-modal data is prevalent. Future research will further explore RNA's capabilities and adaptability to diverse domains and modalities, addressing challenges in unbalanced data distributions, and investigating the integration with other techniques for domain shift mitigation and improved generalization.

Acknowledgements This work was supported in part by the CINI Consortium through the VIDESEC project, in part by FAIR - Future Artificial Intelligence Research, and in part by European Union Next-Generation EU through PIANO NAZIONALE DI RIPRESA E RESILIENZA (PNRR)—MISSIONE 4 COMPONENTE 2, INVESTIMENTO 1.3—D.D. 1555 11/10/2022, under Grant PE00000013. We acknowledge the CINECA award under the IS CRA initiative and IIT HPC infrastructure, for the availability of high performance computing resources and support. The authors would like to thank Prof. Olga Dal Monte for useful discussions on human multi-modal perception.

Funding Open access funding provided by Politecnico di Torino within the CRUI-CARE Agreement.

Data Availability The datasets generated during and/or analysed during the current study are available in the following repositories: EK55: <https://data.bris.ac.uk/data/dataset/3h91syskeag572hl6tvuovwv4d> EK100: <http://dx.doi.org/10.5523/bris.2g1n6qdydwa9u22shpxzpq0t8m> UCF: <https://www.crcv.ucf.edu/data/UCF101.php> HMDB: <https://serrelab.clps.brown.edu/resource/hmdb-a-large-human-motion-database/> ROD: <http://tiny.cc/NRODDatasetDownload> CogBeacon: https://github.com/MikeMpapa/CogBeacon-MultiModal_Dataset_for_Cognitive_Fatigue

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

Aakerberg, A., Nasrollahi, K., & Heder, T. (2017). Improving a deep learning based RGB-D object recognition model by ensemble learning. In *2017 Seventh international conference on image processing theory, tools and applications (IPTA)* (pp. 1–6). IEEE.

- Agarwal, N., Chen, Y.-T., Dariush, B., & Yang, M.-H. (2020). Unsupervised domain adaptation for spatio-temporal action localization. [arXiv:2010.09211](https://arxiv.org/abs/2010.09211).
- Balaji, Y., Sankaranarayanan, S., & Chellappa, R. (2018). Metareg: Towards domain generalization using meta-regularization. In *NeurIPS*.
- Baltrušaitis, T., Ahuja, C., & Morency, L.-P. (2019). Multimodal machine learning: A survey and taxonomy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(2), 423–443.
- Barbato, F., Toldo, M., Michieli, U. & Zanuttigh, P. (2021). Latent space regularization for unsupervised domain adaptation in semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR) workshops* (pp. 2835–2845).
- Bertelson, P. & Gelder, B. D. (2004). The psychology of multimodal perception. In *Crossmodal space and crossmodal attention* (pp. 141–177). Oxford University Press.
- Bo, L., Ren, X., & Fox, D. (2011). Depth kernel descriptors for object recognition. In *2011 IEEE/RSJ international conference on intelligent robots and systems* (pp. 821–826). IEEE.
- Bridle, J., Heading, A., & MacKay, D. (1991). Unsupervised classifiers, mutual information and 'phantom targets. In *Advances in neural information processing systems*, vol. 4. Morgan-Kaufmann.
- Bucci, S., D'Innocente, A., Liao, Y., Carlucci, F. M., Caputo, B., & Tommasi, T. (2021). Self-supervised learning across domains. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(9), 5516–5528.
- Buda, M., Maki, A., & Mazurowski, M. A. (2018). A systematic study of the class imbalance problem in convolutional neural networks. *Neural Network*, 106, 249–259.
- Cannici, M., Plizzari, C., Planamente, M., Ciccone, M., Bottino, A., Caputo, B., & Matteucci, M. (2021). N-rod: A neuromorphic dataset for synthetic-to-real domain adaptation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 1342–1347).
- Carreira, J., & Zisserman, A. (2017). Quo vadis, action recognition? A new model and the kinetics dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 6299–6308).
- Chen, M.-H., Kira, Z., AlRegib, G., Yoo, J., Chen, R., & Zheng, J. (2019). Temporal attentive alignment for large-scale video domain adaptation. In *Proceedings of the IEEE international conference on computer vision* (pp. 6321–6330).
- Chen, M.-H., Li, B., Bao, Y., AlRegib, G., & Kira, Z. (2020). Action segmentation with joint self-supervised temporal domain adaptation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 9454–9463).
- Chen, C., Li, J., Han, X., Liu, X. & Yu, Y. (2022a). Compound domain generalization via meta-knowledge encoding. In *CVPR*.
- Chen, C.-F. R., Panda, R., Ramakrishnan, K., Feris, R., Cohn, J., Oliva, A., & Fan, Q. (2021). Deep analysis of CNN-based spatio-temporal representations for action recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 6165–6175).
- Chen, C., Tang, L., Liu, F., Zhao, G., Huang, Y. & Yu, Y. (2022b). Mix and reason: Reasoning over semantic topology with data mixing for domain generalization. In *NeurIPS*.
- Choi, J., Sharma, G., Chandraker, M., & Huang, J.-B. (2020a). Unsupervised and semi-supervised domain adaptation for action recognition from drones. In *The IEEE winter conference on applications of computer vision* (pp. 1717–1726).
- Choi, J., Sharma, G., Schuler, S., & Huang, J.-B. (2020b). Shuffle and attend: Video domain adaptation. In *European conference on computer vision* (pp. 678–695). Springer.
- Damen, D., Doughty, H., Farinella, G. M., Fidler, S., Furnari, A., Kazakos, E., Moltisanti, D., Munro, J., Perrett, T., Price, W. & Wray,

- M. (2018). Scaling egocentric vision: The epic-kitchens dataset. In *Proceedings of the European conference on computer vision (ECCV)* (pp. 720–736).
- Damen, D., Doughty, H., Farinella, G. M., Furnari, A., Ma, J., Kazakos, E., Moltisanti, D., Munro, J., Perrett, T., Price, W., & Wray, M. (2022). Rescaling egocentric vision: Collection, pipeline and challenges for epic-kitchens-100. *IJCV*.
- Damen, D., Kazakos, E., Price, W., Ma, J., & Doughty, H. (2020). Epic-kitchens-55-2020 challenges report.
- Deng, J., Dong, W., Socher, R., Li, L., Li, K., & Fei-Fei, L. (2009). Imagenet: A large-scale hierarchical image database. In *IEEE conference on computer vision and pattern recognition* (pp. 248–255).
- Deng, Z., Luo, Y., & Zhu, J. (2019). Cluster alignment with a teacher for unsupervised domain adaptation. In *Proceedings of the IEEE international conference on computer vision* (pp. 9944–9953).
- Dou, Q., Coelho de Castro, D., Kamnitsas, K., & Glocker, B. (2019). Domain generalization via model-agnostic learning of semantic features. *Advances in Neural Information Processing Systems*, 32, 6450–6461.
- Driver, J., & Spence, C. (1998). Attention and the crossmodal construction of space. *Trends in Cognitive Sciences*, 2(7), 254–262.
- Ganin, Y., & Lempitsky, V. (2015). Unsupervised domain adaptation by backpropagation, ser. In *Proceedings of machine learning research* vol. 37 (pp. 1180–1189). Lille: PMLR.
- Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., & Lempitsky, V. (2016). Domain-adversarial training of neural networks. *The Journal of Machine Learning Research*, 17(1), 2030–2096.
- Gong, X., Mohan, S., Dhingra, N., Bazin, J.-C., Li, Y., Wang, Z., & Ranjan, R. (2023). MMG-Ego4D: Multi-modal generalization in egocentric action recognition.
- Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., & Smola, A. (2012). A kernel two-sample test. *Journal of Machine Learning Research*, 13(1), 723–773.
- Guo, Y. & Zhang, L. (2017). One-shot face recognition by promoting underrepresented classes.
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 770–778).
- Huang, Y., Lin, J., Zhou, C., Yang, H., Huang, L. (2022). Modality competition: What makes joint training of multi-modal network fail in deep learning? (Provably). In *Proceedings of the 39th international conference on machine learning, series, proceedings of machine learning research* vol. 162 (pp. 9226–9259). PMLR.
- Ioffe, S. & Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift (pp. 448–456).
- Jamal, A., Namboodiri, V. P., Deodhare, D., & Venkatesh, K. (2018). Deep domain adaptation in action space. In *BMVC*.
- Jaritz, M., Vu, T.-H., de Charette, R., Wirbel, E., & Pérez, P. (2020). xMUDA: Cross-modal unsupervised domain adaptation for 3D semantic segmentation. In *CVPR*.
- Kazakos, E., Nagrani, A., Zisserman, A., Damen, D. (2019). Epic-fusion: Audio-visual temporal binding for egocentric action recognition. In *The IEEE international conference on computer vision (ICCV)*.
- Kim, B., & Kim, J. (2020). Adjusting decision boundary for class imbalanced learning. *IEEE Access*, 8, 81674–81685.
- Kim, D., Tsai, Y.-H., Zhuang, B., Yu, X., Sclaroff, S., Saenko, K. & Chandraker, M. (2021). Learning cross-modal contrastive features for video domain adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision* (pp. 13618–13627).
- Kuehne, H., Jhuang, H., Garrote, E., Poggio, T., & Serre, T. (2011). HMDB: A large video database for human motion recognition. In *International conference on computer vision* (pp. 2556–2563). IEEE.
- Lai, K., Bo, L., Ren, X., & Fox, D. (2011). A large-scale hierarchical multi-view RGB-D object dataset. In *IEEE international conference on robotics and automation* (pp. 1817–1824). IEEE.
- Lange, F., Brückner, C., Knebel, A., Seer, C., & Kopp, B. (2018). Executive dysfunction in Parkinson’s disease: A meta-analysis on the Wisconsin card sorting test literature. *Neuroscience & Biobehavioral Reviews*, 93, 38–56.
- Li, M., Cheung, Y.-M., & Jiang, J. (2022). Feature-balanced loss for long-tailed visual recognition. In *IEEE International Conference on Multimedia and Expo (ICME)* (pp. 1–6).
- Li, H., Jialin Pan, S., Wang, S., & Kot, A. C. (2018). Domain generalization with adversarial feature learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 5400–5409).
- Li, Y., Yang, Y., Zhou, W. & Hospedales, T. (2019). Feature-critic networks for heterogeneous domain generalization. In *ICML*.
- Liu, W., Luo, Z., Cai, Y., Yu, Y., Ke, Y., Junior, J. M., Gonçalves, W. N., & Li, J. (2021). Adversarial unsupervised domain adaptation for 3d semantic segmentation with multi-modal learning. *ISPRS Journal of Photogrammetry and Remote Sensing*, 176, 211–221.
- Li, Y., Wang, N., Shi, J., Hou, X., & Liu, J. (2018). Adaptive batch normalization for practical domain adaptation. *Pattern Recognition*, 80, 109–117.
- Loghmani, M. R., Robbiano, L., Planamente, M., Park, K., Caputo, B., & Vincze, M. (2020). Unsupervised domain adaptation through inter-modal rotation for RGB-D object recognition. *IEEE Robotics and Automation Letters*, 5(4), 6631–6638.
- Long, M., Cao, Y., Wang, J., & Jordan, M. (2015). Learning transferable features with deep adaptation networks. In *International conference on machine learning* (pp. 97–105). PMLR.
- Lv, J., Liu, K., & He, S. (2021). Differentiated learning for multi-modal domain adaptation. In *Proceedings of the 29th ACM international conference on multimedia*. ACM.
- Munro, J. & Damen, D. (2020). Multi-modal domain adaptation for fine-grained action recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)*.
- O’Callaghan, C. (2012). *Perception and multimodality*. Oxford: Oxford University Press.
- Pan, B., Cao, Z., Adeli, E., & Niebles, J. C. (2020). Adversarial cross-domain action recognition with co-attention. In *Proceedings of the AAAI conference on artificial intelligence* vol. 34, no. 07 (pp. 11815–11822).
- Papakostas, M., Rajavenkatanarayanan, A., & Makedon, F. (2019). Cog-beacon: A multi-modal dataset and data-collection platform for modeling cognitive fatigue. *Technologies*, 7(2), 46.
- Peng, X., Huang, Z., Lv, J., Zhu, H., & Zhou, J. T. (2019). COMIC: Multi-view clustering without parameter selection. In *Proceedings of the 36th international conference on machine learning, series proceedings of machine learning research* vol. 97 (pp. 5092–5101). PMLR.
- Peng, X., Wei, Y., Deng, A., Wang, D., Hu, D. (2022). Balanced multi-modal learning via on-the-fly gradient modulation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 8238–8247).
- Planamente, M., Goletto, G., Trivigno, G., Averta, G., & Caputo, B. (2022). PoliTO-IIT-CINI submission to the EPIC-KITCHENS-100 unsupervised domain adaptation challenge for action recognition. [arXiv:2209.04525](https://arxiv.org/abs/2209.04525)
- Planamente, M., Plizzari, C., Alberti, E., & Caputo, B. (2022). Domain generalization through audio-visual relative norm alignment in first person action recognition. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision (WACV)* (pp. 1807–1818).
- Plizzari, C., Planamente, M., Alberti, E., & Caputo, B. (2021). PoliTO-IIT submission to the EPIC-KITCHENS-100 unsu-

- pervised domain adaptation challenge for action recognition. [arXiv:2107.00337](https://arxiv.org/abs/2107.00337)
- Qi, F., Yang, X., & Xu, C. (2018). A unified framework for multimodal domain adaptation. In *Proceedings of the 26th ACM international conference on multimedia*. ACM.
- Rahate, A., Walambe, R., Ramanna, S., & Kotecha, K. (May2022). Multimodal co-learning: Challenges, applications with datasets, recent advances and future directions. *Information Fusion*, 81, 203–239.
- Ranjan, R., Castillo, C. D., Chellappa, R. (2017). L2-constrained softmax loss for discriminative face verification. [arXiv:1703.09507](https://arxiv.org/abs/1703.09507)
- Razzaghi, P., Abbasi, K., Shirazi, M., & Shabani, N. (2021). Modality adaptation in multimodal data. *Expert Systems with Applications*, 179, 115126.
- Rebecq, H., Gehrig, D., & Scaramuzza, D. (2018). ESIM: An open event camera simulator. In *Conference on robot learning* (pp. 969–982). PMLR.
- Sahoo, A., Shah, R., Panda, R., Saenko, K., & Das, A. (2021). Contrast and mix: Temporal contrastive video domain adaptation with background mixing. In *Thirty-fifth conference on neural information processing systems*.
- Saito, K., Watanabe, K., Ushiku, Y. & Harada, T. (2018). Maximum classifier discrepancy for unsupervised domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 3723–3732).
- Sevilla-Lara, L., Liao, Y., Güneş, F., Jampani, V., Geiger, A. & Black, M. J. (2019). On the integration of optical flow and action recognition. In *German conference on pattern recognition* (pp. 281–297). Springer.
- Song, X., Zhao, S., Yang, J., Yue, H., Xu, P., Hu, R., & Chai, H. (2021). Spatio-temporal contrastive domain adaptation for action recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 9787–9795).
- Soomro, K., Zamir, A. R., & Shah, M. (2012). UCF101: A dataset of 101 human actions classes from videos in the wild. [arXiv:1212.0402](https://arxiv.org/abs/1212.0402).
- Stein, B. E., Laurienti, P. J., Wallace, M. T. & Stanford, T. R. (2002). Multisensory integration. In *Encyclopedia of the human brain* (pp. 227–241). Elsevier.
- Sun, B. & Saenko, K. (2016). Deep coral: Correlation alignment for deep domain adaptation. In *ECCV*.
- Sun, S., Kuang, Z., Sheng, L., Ouyang, W., & Zhang, W. (2018). Optical flow guided feature: A fast and robust motion representation for video action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 1390–1399).
- Tang, H., & Jia, K. (2020). Discriminative adversarial domain adaptation. In *AAAI* (pp. 5940–5947).
- Volpi, R., Namkoong, H., Sener, O., Duchi, J. C., Murino, V., & Savarese, S. (2018). Generalizing to unseen domains via adversarial data augmentation. In *Advances in neural information processing systems* (pp. 5334–5344).
- Wang, Y., Li, H. & Kot, A. C. (2020). Heterogeneous domain generalization via domain mixup. In *ICASSP*.
- Wang, X., Li, L., Ye, W., Long, M., & Wang, J. (2019). Transferable attention for domain adaptation. In *Proceedings of the AAAI conference on artificial intelligence* vol. 33 (pp. 5345–5352).
- Wang, W., Tran, D., & Feiszli, M. (2020). What makes training multi-modal classification networks hard?. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 12695–12705).
- Wang, F., Xiang, X., Cheng, J., & Yuille, A. L. (2017). Normface: L2 hypersphere embedding for face verification. In *Proceedings of the 25th ACM international conference on multimedia* (pp. 1041–1049).
- Wang, L., Xiong, Y., Wang, Z., Qiao, Y., Lin, D., Tang, X., & Van Gool, L. (2016). Temporal segment networks: Towards good practices for deep action recognition. In *European conference on computer vision* (pp. 20–36). Springer.
- Wei, P., Kong, L., Qu, X., Yin, X., Xu, Z., Jiang, J., & Ma, Z. (2022). Unsupervised video domain adaptation: A disentanglement perspective. [arXiv:2208.07365](https://arxiv.org/abs/2208.07365).
- Wu, Y., Liu, H., Li, J., Fu, Y. (2017). Deep face recognition with center invariant loss. In *Proceedings of the on thematic workshops of ACM multimedia 2017, series thematic workshops'17* (pp. 408–414). New York: Association for Computing Machinery.
- Wu, X., Zhang, S., Zhou, Q., Yang, Z., Zhao, Z., & Latecki, L. J. (2021). Entropy minimization versus diversity maximization for domain adaptation. *IEEE Transactions on Neural Networks and Learning Systems*, 1–12.
- Xu, R., Li, G., Yang, J., & Lin, L. (2019). Larger norm more transferable: An adaptive feature norm approach for unsupervised domain adaptation. In *2019 IEEE/CVF international conference on computer vision (ICCV)* (pp. 1426–1435). Los Alamitos: IEEE Computer Society.
- Xu, R., Li, G., Yang, J., & Lin, L. (2019). Larger norm more transferable: An adaptive feature norm approach for unsupervised domain adaptation. In *Proceedings of the IEEE international conference on computer vision* (pp. 1426–1435).
- Xu, M., Zhang, J., Ni, B., Li, T., Wang, C., Tian, Q. & Zhang, W. (2020). Adversarial domain adaptation with domain mixup. In *AAAI*.
- Yang, L., Huang, Y., Sugano, Y., & Sato, Y. (2022). Interact before align: Leveraging cross-modal knowledge for domain adaptive action recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 14722–14732).
- Yang, M., Li, Y., Hu, P., Bai, J., Lv, J., & Peng, X. (2022). Robust multi-view clustering with incomplete information. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(1), 1055–1069.
- Yang, J., Qian, H., Zou, H., & Xie, L. (2021). Learning decomposed hierarchical feature for better transferability of deep models. *Information Sciences*, 580, 385–397.
- Yao, Z., Wang, Y., Wang, J., Yu, P., & Long, M. (2021). Videodg: Generalizing temporal relations in videos to novel domains. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44, 7989–8004.
- Zhang, Y., Li, M., Li, R., Jia, K., & Zhang, L. (2022). Exact feature distribution matching for arbitrary style transfer and domain generalization. In *CVPR*.
- Zheng, Y., Pal, D. K., & Savvides, M. (2018). Ring loss: Convex feature normalization for face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 5089–5097).
- Zhou, B., Andonian, A., Oliva, A., & Torralba, A. (2018). Temporal relational reasoning in videos. In *Proceedings of the European Conference on Computer Vision (ECCV)* (pp. 803–818).
- Zhou, Q., Zhou, W., Wang, S., & Xing, Y. (2020). Unsupervised domain adaptation with adversarial distribution adaptation network. *Neural Computing and Applications*, 33(13), 7709–7721.
- Zhu, H., Luo, M.-D., Wang, R., Zheng, A.-H., & He, R. (2021). Deep audio-visual learning: A survey. *International Journal of Automation and Computing*, 18(3), 351–376.
- Zisserman, A., Carreira, J., Simonyan, K., Kay, W., Zhang, B., Hillier, C., Vijayanarasimhan, S., Viola, F., Green, T., Back, T., Natsev, P., & Suleyman, M. (2017). The kinetics human action video dataset.
- Zou, Y., Yu, Z., Liu, X., Kumar, B., & Wang, J. (2019). Confidence regularized self-training. In *Proceedings of the IEEE/CVF international conference on computer vision* (pp. 5982–5991).