

# Learning with privileged information via adversarial discriminative modality distillation

Nuno C. Garcia, Pietro Morerio, and Vittorio Murino, *Senior Member, IEEE*

**Abstract**—Heterogeneous data modalities can provide complementary cues for several tasks, usually leading to more robust algorithms and better performance. However, while training data can be accurately collected to include a variety of sensory modalities, it is often the case that not all of them are available in real life (testing) scenarios, where a model has to be deployed. This raises the challenge of how to extract information from multimodal data in the training stage, in a form that can be exploited at test time, considering limitations such as noisy or missing modalities. This paper presents a new approach in this direction for RGB-D vision tasks, developed within the adversarial learning and privileged information frameworks. We consider the practical case of learning representations from depth and RGB videos, while relying only on RGB data at test time. We propose a new approach to train a hallucination network that learns to distill depth information via adversarial learning, resulting in a clean approach without several losses to balance or hyperparameters. We report state-of-the-art results on object classification on the NYUD dataset and video action recognition on the largest multimodal dataset available for this task, the NTU RGB+D, as well as on the Northwestern-UCLA.

**Index Terms**—Multimodal deep learning, adversarial learning, privileged information, network distillation, modality hallucination.



## 1 INTRODUCTION

DEPTH perception is the ability to reason about the 3D world, critical for the survival of many hunting predators and an important skill for humans to understand and interact with the surrounding environment. It develops very early in humans when babies start to crawl [1], and emerges from a variety of mechanisms that jointly contribute to the sense of relative and absolute position of objects, called depth cues. Besides binocular cues (*e.g.* stereovision), humans use monocular cues that relate to *a priori* visual assumptions derived from 2D single images through shadows, perspective, texture gradient, and other signals (*e.g.* the assumption that objects look blurrier the further they are, or that if an object occludes another it must be closer, *etc*) [2]. As matter of fact, although object distance was observed to be underestimated in human monocular vision [3], we are still able to perform most of our vision-related tasks with good efficiency, even with one eye covered.

Similarly, depth perception is often of paramount importance for many computer vision tasks related to robotics, autonomous driving, scene understanding, to name a few. The emergence of cheap depth sensors and the need for big data led to big multimodal datasets containing RGB, depth, infrared, and joints sequences [4], which in turn stimulated multimodal deep learning approaches. Traditional computer vision tasks like action recognition, object detection, or instance segmentation have been shown to benefit performance gains if the model considers other modalities, namely depth, instead of RGB only [5], [6], [7], [8].

However, it is reasonable to expect that depth data is not going to be always available when a model is deployed




Fig. 1. What is the best way of using all data available at training time, considering a missing (or noisy) modality at test time?

in a real scenario, either due to the impossibility to collect depth data with enough quality (*e.g.* due to far-distance or reflectance issues) or to install depth sensors everywhere, sensor or communications failure, or other unpredictable events.

Considering this limitation, we would like to answer the following question: what is the best way of using all data available at training time, in order to learn robust representations, knowing that there are missing (or noisy) modalities at test time? In other words, is there any added value in training a model by exploiting multimodal data, even if only one modality is available at test time? Unsurprisingly, the simplest and most commonly adopted solution consists in training the model using only the modality in which it will be tested. However, a more interesting alternative

N. C. Garcia, P. Morerio and V. Murino are with Pattern Analysis & Computer Vision (PAVIS), Istituto Italiano di Tecnologia (IIT), Genova, Italy.  
N. C. Garcia is also with University of Genova, Italy.  
V. Murino is also with Dept. of Computer Science, University of Verona, Italy.  
Primary email contact: nuno.garcia@iit.it.

is trying to exploit the potential of the available data and train the model using all modalities, being however aware of the fact that not all of them will be accessible at test time. This learning paradigm, i.e., when the model is trained using extra information, is generally known as *learning with privileged information* [9] or *learning with side information* [10].

In this work, we propose an adversarial discriminative modality distillation (ADMD) strategy within a multimodal-stream framework that learns from different data modalities and can be deployed and tested on a subset of these. More in detail, our model learns from RGB *and* depth video sequences and is tested on RGB only, but due to its general design it can also be used to manage whatever combination of other modalities as well. We evaluate its performance on the task of video action recognition and object classification. To this end, we introduce a new adversarial learning strategy to learn a hallucination network (Fig. 2), whose goal is to mimic the test time missing modality features, while preserving their discriminative power. The hallucination network uses RGB only as input and tries to recover useful depth features for the task at hand. We would like to stress the fact that, in contrast to estimating real depth maps from RGB, we operate at feature level.

Such net can be interpreted as a source of the aforementioned monocular depth cues, i.e. a source of depth cues from a single 2D RGB image.

On the one hand, our work is inspired by previous works using hallucination networks in the context of learning with privileged information. This was primarily proposed in [10], which presented an end-to-end single step training method to learn a hallucination network. This work was recently revisited in [11] considering a multi-step learning paradigm using a loss inspired in the generalized distillation framework [12]. On the other hand, adversarial learning has been shown to be a powerful tool to model data distributions [13], [14]. Building upon these ideas, we propose a new approach to learn the hallucination network via a discriminative adversarial learning strategy. Our proposed method has several advantages: it is agnostic regarding the pair of modalities used, which greatly simplifies its extension beyond RGB and depth data, and it is able to deal with videos by design, by exploiting the time as auxiliary information. Furthermore, it dumps the need to balance the different losses used in the other methods [10] [11]. Finally, thanks to the discriminator design, which includes an auxiliary classification task, our method is able to fully transfer the discriminative capability from the teacher network to the student stream, up to a full recovery of the teacher’s accuracy.

To summarize, the main contributions of this paper are the following:

- we propose a new, flexible, and better performing approach to learn a hallucination stream within a multimodal-stream network architecture, which consists in an adversarial learning strategy that effectively exploits multiple data modalities at training while using only one at test time;
- more technically, we propose a discriminator network which is time-aware, and jointly solves 1) the classical binary classification task (real/generated)

and 2) an auxiliary task, which inherently endows the learned features with discriminative power;

- we report results – in the privileged information scenario – on the NYUD dataset for the task of object classification and on the large-scale NTU RGB+D dataset [15] and the Northwestern-UCLA [16] for the task of action recognition.

The rest of the paper is organized as follows. Section 2 relates this work to literature in adversarial learning, privileged information, and multimodal deep learning. Section 3 presents the details of the proposed architecture and the novel learning strategy. Section 4 reports the results on object recognition and video action recognition datasets, while comparing them to the current state of the art. It also investigates how the different parts of our approach contribute to the performance of the method through an ablation study conducted on a subset the NTU RGB+D dataset. Finally, we draw conclusions and future research directions in section 5.

## 2 RELATED WORK

Our work is at the intersection of four topics: adversarial learning [13], RGB-D vision, network distillation [17] and privileged information [9]. However, as Lopez *et al.* noted, privileged information and network distillation are instances of the same more inclusive theory, called generalized distillation [12].

### 2.1 Generalized Distillation

Within the generalized distillation framework, our model is both related to the privileged information theory [9], considering that the extra modality (depth, in this case) is only used at training time, and to the distillation framework, considering that our hallucination network is effectively learning by distilling the knowledge of a the previously learned “teacher” network, despite not using a distillation loss.

In this context, the closest works to our approach are [10] by Hoffman *et al.* and [11] by Garcia *et al.*

The work of Hoffman *et al.* [10] introduced a model to hallucinate depth features from RGB input for object detection task. While the idea of using a hallucination stream is similar to the one thereby presented, the mechanism used to learn it is different. In [10], the authors use an Euclidean loss between the depth and hallucinated feature maps, that is part of the total loss along with more than ten classification and localization losses. This makes its effectiveness dependent on hyperparameter tuning to balance the different values, as the model is trained jointly in one step by optimizing the aforementioned composite loss.

In [11], Garcia *et al.* built on this idea to propose a new staged training procedure, that would effectively lead to learn a better teacher network, and a new loss inspired to the distillation framework, that not only uses the Euclidean distance between feature maps, and the one-hot labels, but also leverages soft predictions from the teacher network (the depth stream, in this case). Moreover, the authors encouraged the learning by design, by using multiplier cross-stream connections [18], and extended the method to

a video task, namely video action recognition. Differently from [11] and [10], we propose an adversarial strategy to learn the hallucination stream, which alleviates the need for balancing losses or tuning hyperparameters.

An interesting work lying at the intersection of multi-modal learning and learning with privileged information is ModDrop by Neverova *et al.* [19]. Here the authors propose a modality-based dropout strategy, where each input modality is *entirely* dropped (actually zeroed) with some probability during training. The resulting model is proved to be more resilient to missing modalities at test time. We compare with ModDrop in the task of object classification.

Luo *et al.* [20] addressed a similar problem, where the model is first trained on several modalities (RGB, depth, joints and infrared), but tested only in one. The authors propose a graph-based distillation method that is able to distill information from all modalities at training time, while also passing through a validation phase on a subset of modalities. This approach achieves state-of-the-art results in action recognition and action detection tasks. Our work substantially differs from [20] since we benefit from a hallucination mechanism, consisting in an auxiliary hallucination network trained by leveraging a previously trained network. This mechanism allows the model to learn to emulate the presence of the missing modality at test time.

Finally, another recent related work is [21], which proposes a distillation framework where there is no frozen teacher network, but all the networks work as an ensemble that learn in a collaboratively manner.

## 2.2 RGB-D vision

Video action recognition and object detection have a long and rich field of literature, spanning from classification methods using handcrafted features, *e.g.* [22], [23], [24], [25], [26], [27] to modern deep learning approaches, *e.g.* [7], [28], [29], [30], [31], using either RGB-only or together with depth data. We point to some of the more relevant works in video action recognition and object recognition using RGB and depth, including state-of-the-art methods considering the NTU RGB+D and the NYU-Depth V2 datasets, as well as architectures related to our proposed model.

### 2.2.1 Video action recognition

The two-stream model introduced by Simonyan and Zisserman [32] is a landmark on video analysis, and since then has inspired a series of variants that achieved state-of-the-art performance on diverse datasets. This architecture is composed by an RGB and an optical flow stream, which are trained separately, and then fused at the prediction layer. Our model relates to this since the test-time predictions result from the average of the hallucination stream and the RGB stream logits. In [18], the authors propose a variation of the latter, which models spatiotemporal features by injecting the motion stream’s signal into the residual unit of the appearance stream. They also employ 1D temporal convolutions along with 2D spatial convolutions, which we also adopt in this paper. Indeed, the combination of 2D spatial and 1D temporal convolutions has shown to learn better spatiotemporal features than 3D convolutions [33]. The current state of the art in video action recognition [34]

uses 3D temporal convolutions and a new building block dedicated to capture long range dependencies, using RGB data only.

Instead, in [5] the authors explore the complementary properties of RGB and depth data, taking the NTU RGB+D dataset as testbed. They propose a deep autoencoder architecture and a structured sparsity learning machine, and show to achieve state-of-the-art results for action recognition. Liu *et al.* [6] also use RGB and depth complementary information to devise a method for viewpoint invariant action recognition. Here, dense trajectories from RGB data are first extracted, which are then encoded in viewpoint invariant deep features. The RGB and depth features are then used as a dictionary to predict the test label.

To the best of our knowledge, these are state-of-the-art approaches to exploit RGB+D for video action recognition, that report results on the NTU RGB+D dataset [15], the largest video action recognition dataset to offer RGB and depth. However, it is important to note that we propose a fully convolutional model that exploits RGB and depth data at training time only, and uses exclusively RGB data as input at test time. This work goes in the direction of closing the performance gap between privileged information and traditional approaches.

### 2.2.2 Object recognition

Over the years, object recognition based on RGB and depth have been an insightful task to reason on the complementarity of these two modalities, and how depth data should be handled differently or not, compared to RGB. An example of this is [7], in which the authors propose to encode depth images using a geocentric embedding that encodes height above ground and angle with gravity for each pixel in addition to the horizontal disparity, showing that it works better than using raw depth. Differently, in [31], the authors focus on carefully designing a convolutional neural network including a multimodal layer to fuse RGB and depth. Our work differs from these approaches since we focus on learning a model that has access to depth only at training time, which fundamentally changes the feature learning approach.

## 2.3 Adversarial Learning

In the seminal paper of Goodfellow *et al.* [13], the authors propose a generative model that is trained by having two networks playing the so called minimax game. A generator network is trained to generate images from noise vectors, and a discriminator network is trained to classify the generated images as false, and images sampled from the dataset as true. As the game evolves, the generator becomes better and better at generating samples that look like the true images from the data distribution.

Many papers extended this approach in different directions, such as disentangling semantic concepts [35], network compression [36], feature augmentation [37], image to image translation [38], and explored different losses [39] and other tricks to improve performance and stability [40] [41]. Our work relates to this body of work, as the hallucination network of our model tries to generate features in the missing modality feature space through adversarial learning. However, not only the context here is different, since adversarial

learning is explored in the framework of privileged information, but also the task assigned to the discriminator is not the standard one typically used in adversarial learning, as detailed in Section 3.

An important variant of the GAN framework are Conditional GANs (CGANs) [42], that propose to concatenate the label of desired class to be generated, to the noise vector. This mechanism is related, yet different, to how our generator network implicitly conditioned in this work. Our generator network input is a small volume of 5 RGB frames, and temporal convolutions are zero-padded to maintain the volume’s size along the time dimension. The generator is thus implicitly aware of the temporal ordering, since features generated for the first and last frames will heavily be affected by the border effect of zero padding, which is performed several times (at each residual block). To solve this issue, we provide the temporal ordering conditioning label to the discriminator as well.

In a very recent work, Roheda *et al.* also approach the problem of missing modalities in the context of adversarial learning [43]. They address the binary task of person detection using images, seismic, and acoustic sensors, where the latter two are absent at test time. A CGAN is conditioned on the available images and the generator maps a vector noise to representative information from the missing modalities, with an auxiliary L2 loss. In contrast to this work, our CGAN model learns a mapping directly from the test modality to the feature space of the missing modality, with no auxiliary loss. Besides this, we propose a two-step training procedure in order to learn a better teacher network, and provide a stable target for the generator. Finally, we focus on the more demanding tasks of video action recognition and on object recognition.

### 3 LEARNING TO HALLUCINATE DEPTH FEATURES

As stated in the previous sections, our goal is to train a hallucination network that, having as input RGB, is able to produce similar features to the ones produced by the depth network. The reasoning behind this idea is that on one hand depth and RGB provide complementary information for the task, but on the other hand RGB alone contains some cues for depth perception. Therefore, the goal of the hallucination network is to extract such complementary information, that depth data would provide, from RGB frames. We would like to emphasize again that we are interested in recovering useful depth *features*, in contrast to estimating real depth maps from RGB.

This is accomplished in a two-step training procedure, illustrated in Fig. 2, and described in the following. The *first step* (Fig. 2, top) consists in training the RGB and depth streams individually, with the respective input modality, as two standard, separate, supervised learning problems. The resulting ensemble, obtained by fusing the predictions of the two sub-networks (not fine-tuned), represents the full model (TwoStream) that can be exploited when both modalities are available at test time. Its accuracy should be an upper bound for the model we are proposing. In the *second step* (Fig. 2, bottom), we actually train the hallucination network by means of the proposed adversarial learning strategy. As the hallucination network is trained in the context of adversarial

learning, it can be also interpreted as the generator network in the traditional GAN framework [13], since it tries to generate depth features. However, strictly speaking, it is clearly to be considered as an encoder, which tries to extract monocular depth features from RGB input data. The test time setup of step 2 is again a two-stream model (not fine-tuned), composed by the RGB and hallucination networks, both having RGB data as input, assumed to be the only modality available.

#### 3.1 Training procedure

Inspired by the generalized distillation paradigm, we follow a staged learning procedure, where the “teacher” net is trained first (Step 1) and separately from the “student” (Step 2). This is in contrast with [10], where everything is learned end-to-end, but in line with [11], where separating the learning steps is proved to be more effective.

**Step 1.** The RGB and depth streams are trained separately, which is common practice in two-stream architectures. Both depth and appearance streams are trained by minimizing the cross-entropy loss, after being initialized with a pre-trained ImageNet model for all experiments. We test both streams individually and in a two-stream setup, where the final prediction results from the average of the two streams’ logits. We found that fine-tuning the two-stream model does not increase performance consistently. This step can also be regarded as training the teacher network - depth stream - for the next step (see Fig. 2, top).

**Step 2.** The depth stream  $E_d$ , trained in the previous step, is now frozen, in order to provide a stable target for the hallucination network (generator)  $H$ , which plays the adversarial game with a discriminator  $D$  (see Fig. 2, bottom). The primary task of the discriminator  $D$  is to distinguish between the features  $F_H$  generated by the hallucination network  $H$  and  $F_d$  generated by the depth network  $E_d$ . However, as already mentioned, the discriminator is also assigned an auxiliary discriminative task, as detailed in the following.

The networks  $E_d$  and  $H$  architecture are a mix of 2D and 3D convolutions that process a set of frames, and output a feature vector for every frame  $t$  of the input volume, *i.e.*  $F_H^t$  and  $F_d^t$ . This means that each frame can have a different feature vector even if sampled from the same video, depending on its dynamics and its position  $t$  in the input volume. For example, the first frame of the action “shaking hand” might be very different from the middle frame, whereas for “reading” it might be similar. This increases the complexity for the generator, that have not only to generate features similar to  $F_d$ , but also to match the order in which they are generated. Namely,  $F_H^t$  should be similar to  $F_d^t$ , for every frame  $t$  of the input volume. We ease this issue by providing as input to  $D$  the one-hot encoding vector of the index  $t$ , which we denote  $y^t$ , concatenated with the respective feature vector, which relates to the CGAN mechanism [42].

In standard adversarial training, the discriminator  $D$  would try to assign the binary label true/fake to the feature vector coming from the two different streams. However, we found that features  $F_H$  generated with this mechanism, although being very well mixed and indistinguishable from

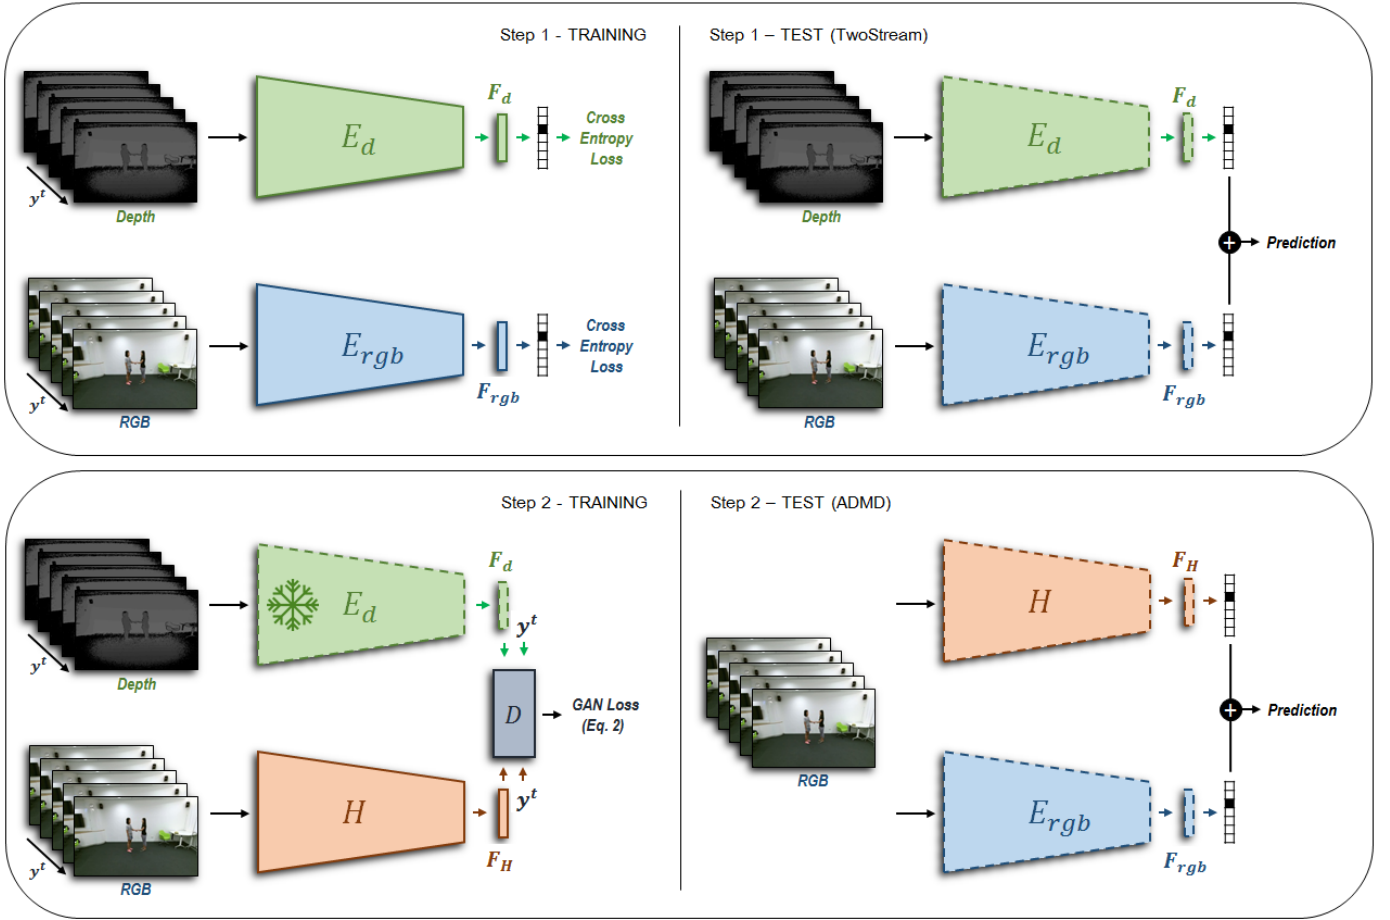


Fig. 2. Architecture and training steps (solid lines - module is *trained*; dashed lines - module is *frozen*). **Step 1:** Separate pretraining of RGB and Depth networks (Resnet-50 backbone with temporal convolutions). The bottleneck described in section 3.2 is highlighted as a separate component. At test time the raw predictions (logits) of the two separate streams are simply averaged. The complementary information carried by the two streams bring a significant boost in the recognition performance. **Step 2:** The depth stream is frozen. The hallucination stream  $H$  is initialized with the depth stream’s weights and adversarially trained against a discriminator. The discriminator is fed with the concatenation of the bottleneck feature vector and the temporal frame ordering label  $y^t$ , as detailed in Section 3.1. The discriminator also features an additional classification task, i.e. not only it is trained to discriminate between hallucinated and depth features, but also to assign samples to the correct class (Eq. 2). The hallucination stream thus learns monocular depth features from the depth stream while maintaining discriminative power. At test time, predictions from the RGB and the hallucination streams are fused.

$F_d$ , were struggling to provide a good accuracy for all the considered tasks, *i.e.* were lacking discriminative power. For this reason we assign to the discriminator the auxiliary task of classifying feature vectors with their correct class.

The adversarial learning problem is formalized as follows. Consider the RGB-D dataset  $(X_{rgb}, X_d, Y)$  where  $x_{rgb}^t, x_d^t \sim (X_{rgb}, X_d)$  are time aligned RGB and depth frames,  $y \sim Y$ , is the  $C$ -dimensional one-hot encoding of the class label, and  $C$  is the number of classes for the problem at hand.

Now, define the *extended label vector* with  $C + 1$  components (classes):

$$\hat{y} = \begin{cases} [0, \dots, 0, 1], & \text{for } x_{rgb} \\ [y_i, 0] & \text{for } x_d \end{cases} \quad (1)$$

Using this label vector instead of the classical 0/1 (real/generated) binary label in the discriminator encourages feature representations  $F_H$  learned by  $H$  to encode not only depth (monocular) features, but also discriminative

ones, either purely RGB or depth. This is possibly why the hallucination network often recovers the accuracy of the teacher and sometimes performs even better, as further explained in the experimental section.

Based on the above definitions, we define the following minimax game:

$$\min_{\theta_D} \max_{\theta_H} \ell = \mathbb{E}_{(x_i, y_i) \sim (X_{rgb}, Y)} \mathcal{L}(D(H(x_i) || y^t), \hat{y}_i) + \mathbb{E}_{(x_i, y_i) \sim (X_d, Y)} \mathcal{L}(|D(E_d(x_i) || y^t), \hat{y}_i) \quad (2)$$

where  $\theta_H$  and  $\theta_D$  indicate the parameters of the hallucination stream  $H$  and of the discriminator  $D$ ,  $||$  denotes a concatenation operation and  $\mathcal{L}$  is the softmax cross-entropy function. Eq. 2 is optimized via the well known “label flipping hack” [44], which makes the loss function easier to minimize in practice.

### 3.2 Architectural details

All three networks (depth stream -  $E_d$ , RGB stream -  $E_{rgb}$ , and hallucination stream  $H$ ) are modified Resnet-50 [45]

augmented with *temporal convolutions* and endowed with a final *bottleneck layer*.

**Temporal convolutions.** 1D temporal convolutions are inserted in the second residual unit of each ResNet layer as illustrated in Fig. 3, following the recent work of Feichtenhofer *et al.* [18]. For layer  $l$ , the weights  $W_l \in \mathbb{R}^{1 \times 1 \times 3 \times C_l \times C_l}$  are convolutional filters initialized as identity mappings at feature level, and centered in time, where  $C_l$  is the number of channels in layer  $l$ . More in detail, all the  $[1 \times 1 \times 3]$  temporal kernels contained in  $W_l$  are initialized as  $[0, 1, 0]$ , *i.e.* only the information of the central frame is used at the beginning. This progressively changes as training goes on. Very recently, in [33], the authors explored various network configurations using temporal convolutions, comparing different combinations for the task of video classification. This work suggests that decoupling 3D convolutions into 2D (spatial) and 1D (temporal) filters is the best setup in action recognition tasks, producing best accuracies. The intuition for the latter setup is that factorizing spatial and temporal convolutions in two consecutive convolutional layers eases training of the spatial and temporal tasks (also in line with [46]).

**Bottleneck.** Generating, encoding, or aligning high dimensional feature vectors via adversarial training is often a difficult task, due to the inherent instability of the saddle point defined by the GAN minimax game. For this reason, [37] proposes to align a lower dimensional vector, obtained by adding a *bottleneck layer* to standard architectures. This usually does not affect performances of baseline models.

Indeed, the size of the last ResNet-50 layer (before the logits) is  $[7, 7, 2048]$ , or simply  $[2048]$  after pooling. For this reason, we further modify the ResNet-50 by adding either i) an additional convolutional layer, whose weights  $W_b \in \mathbb{R}^{7 \times 7 \times 2048 \times 128}$ , applied with no padding, reduce the dimensionality to 128; or ii) a simple 128-dim fully connected layer after pooling. In Section 4.2.1 we further explore the choice of the bottleneck.

**Inputs.** For the task of action recognition, the input to the stream networks is made of 5 3-channel frames, uniformly sampled from each video sequence, which motivates temporal convolution. Instead, for the task of object classification (from single images), no temporal kernels are added to the architecture. We try different encodings for the depth channel: for the task of action recognition they are encoded into color images using a jet colormap, as in [47]; for the object recognition task, HHA encoding [48] is already provided in the dataset considered.

**Discriminator.** The discriminator used to play the adversarial game has different architectures depending on the




Fig. 3. Detail of the ResNet residual unit with temporal convolutions (blue block).

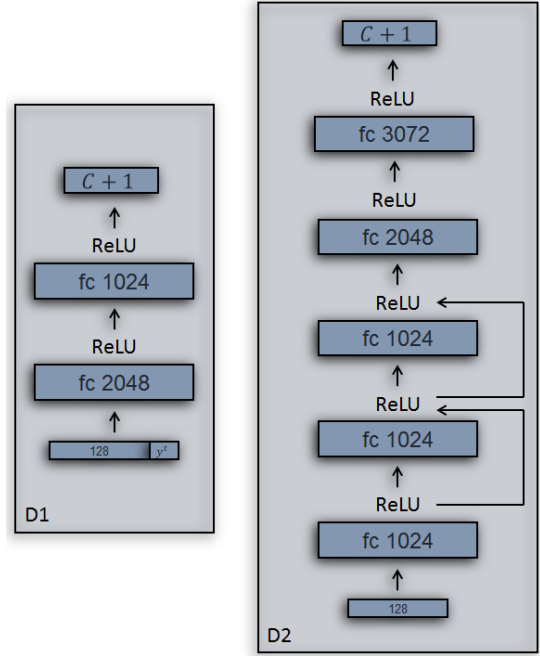


Fig. 4. Architectures for the discriminators used for the two different tasks. Left: D1, for object recognition. Right: D2 for action recognition.

task. Its basic structure is that of a multilayer perceptron, stacking fully connected (fc) layers only, since it takes a vector as input (bottleneck features, possibly concatenated with temporal ordering for tasks involving time). For the task of action recognition, the structure is quite shallow, consisting in  $D1=[fc(2048), fc(1024), fc(C+1)]$ . For the task of object classification the structure is instead more complex  $D2=[fc(1024), fc(1024), fc(1024), fc(2048), fc(3072), fc(C+1)]$ , with skip connections in the lower layers. Being the former discriminator quite deep, residual connections were inserted in order to allow gradient to flow through the underlying hallucination stream. Details of the architectures are sketched in Fig. 4.

## 4 EXPERIMENTS

### 4.1 Datasets

We evaluate the performance of our method on one object classification and two video action classification datasets. For both tasks the model is initialized with ImageNet pre-trained weights. For the experiments on the smaller action recognition dataset NWUCLA, we fine-tune the model starting from the RGB and depth streams trained on the bigger NTU RGB+D dataset.

**NTU RGB+D [15].** This is the largest public dataset for multimodal video action recognition. It is composed by 56,880 videos, available in four modalities: RGB videos, depth sequences, infrared frames, and 3D skeleton data of 25 joints (RGB and depth examples illustrated in Fig. 1). It was acquired with a Kinect v2 sensor in 80 different viewpoints, and includes 40 subjects performing 60 distinct actions. We follow the two evaluation protocols originally proposed in [15], which are cross-subject and cross-view. As in the original paper, we use about 5% of the training data as validation set for both protocols. The masked depth maps



Fig. 5. Examples of RGB and depth frames from the NYUD (RGB-D) dataset.

are converted to a three channel map via a jet mapping, as in [47].

**Northwestern-UCLA [16].** Similarly to the other datasets, it provides RGB, depth and skeleton sequences for 1475 samples. It features 10 subjects performing 10 actions captured in 3 different views.

**NYUD (RGB-D).** The is a dataset of objects (see examples in Fig. 5 gathered by cropping out tight bounding boxes around instances of 19 object classes present in the NYUD [49] dataset. It comprises 2,186 paired labeled training images and 2,401 test images (RGB-D). Depth images are HHA-encoded [48]. The dataset was proposed in [10] but also employed in [37], [50], [51] for the task of modality adaptation, in the framework of domain adaptation (train on one modality, adapt and test the model on the other modality). The task here is object classification, training on both modalities and testing on RGB only.

## 4.2 Ablation Study

The ablation study is performed on a part of the NTU RGB+D dataset, designated as mini-NTU, which consists in random samples from the training set, considering approximately a third of the original dataset size. The test set is the same used in the other experiments [15].

We study how the hallucination network performance is affected by (1) feeding different types of input to the discriminator, and (2) having the discriminator to perform different tasks.

### 4.2.1 Bottleneck size

The discriminator receives as input the feature vector  $F_H$  or  $F_d$  from either the hallucination stream or the depth stream, respectively, along with the frame index label  $y^t$ . It is known that a too big feature vector may cause the GAN training to underperform [37], which we also observe in our experiments, reported in Table 1.

We first trained our depth network without bottleneck on the full NTU dataset, reaching 70.53% accuracy. This network is then used as target to learn the hallucination model. We observed that the hallucination model trained without bottleneck, *i.e.*, the input to the discriminator is the 2048-dimensional feature vector, is far from recovering the performance of the target (reaching only 54.25%), even if the training space is reduced to the NTU-mini dataset (60.95%).

TABLE 1  
Ablation Study - Bottleneck size. Hallucination network underperforming with  $F_x \in \mathbb{R}^{2048}$ .

Network	Dataset	X-Subject
Depth stream, normal - (target)	NTU	70.53%
Hall. net, $F_x \in \mathbb{R}^{2048}$	NTU	54.25%
Hall. net, $F_x \in \mathbb{R}^{2048}$	NTU-mini	60.95%
Depth stream, w/ bottleneck - (target)	NTU	69.13%
Hall. net, $F_x \in \mathbb{R}^{128}$	NTU	72.14%

TABLE 2  
Ablation Study - Investigating different bottleneck implementations. The Table reports Hallucination network performances on NTU-mini.

Depth stream - versions	X-Subject	X-View
Depth stream wo/ bottleneck	63.95%	62.70%
One conv	55.64%	57.91%
Spatial conv + 1D conv	53.21%	52.58%
pool + conv	61.41%	63.15%

We then train a network with a 128-dimensional bottleneck (69.13%), initialized with the previous depth stream, except for the bottleneck that is randomly initialized with the MSRA initialization [52]. The hallucination model that learns using the bottleneck feature vector is able not only to recover, but to surpass the performance of the depth stream, reaching 72.14% accuracy. We observed this behaviour in other experiments along the paper, and we comment that later in this section (Sect. 4.4).

### 4.2.2 Bottleneck implementation

In Table 2 we investigate different ways to decrease the size of  $F_x$  from  $\mathbb{R}^{2048}$  to  $\mathbb{R}^{128}$ , as suggested in [37]. After the last feature map, which is of dimension  $7*7*2048$ , we tested the three following ways:

- convolution of [128,7,7] to  $1*1*128$ ,
- spatial convolution of [7,7] to  $1*1*2048$  followed by 1D convolution to  $1*1*128$ , and
- pooling layer to  $1*1*2048$  followed by 1D convolution to  $1*1*128$

Even though the depth stream is just trained on the NTU-mini (63.95% for cross subject, and 62.70% for cross view), the hallucination stream that implements the pool+conv bottleneck is able to recover almost completely (61.41% for cross subject), or even surpass (63.15% for cross view), the original depth stream performance. This was the architectural choice we used in the rest of the experiments.

### 4.2.3 Discriminator: inputs and tasks

In this section, we explore whether the task assigned to the discriminator influences the hallucination performance. As introduced in Section 3, our hypothesis is that the generator has the difficult task of generating frames that not only correspond to depth frames, but also need to be paired with these. We solve this by introducing the additional information of the frame index  $y^t$ , which explicits the desired alignment. Table 3 shows results regarding the (1) traditional binary task of a GAN generator having as

TABLE 3  
Ablation Study - Investigating different inputs and tasks for the discriminator. The Table reports Hallucination network performances (NTU-mini).

Input	Task	X-Subject
Teacher network (pool + conv, Table 2)	-	61.41%
F(x)	0/1 classification	1.81%
F(x)	$\hat{y}$ classification	59.87%
F(x)    $y_t$	$\hat{y}$ classification	63.03%

input the feature bottleneck, (2) the  $\hat{y}$  classification task having the same input as before, and (3) the proposed approach. The traditional binary task converges to a perfect equilibrium, but the hallucination stream’s accuracy is close to random chance, meaning that the learned features are not discriminant at all. While the second approach is able to learn, the inclusion of the frame order supervision  $y^t$  shows to increase the performance. It is reasonable that this mechanism produces maximized gains on more challenging and diverse datasets, as the complete NTU dataset, or in fully 3d-convolutional architectures such as I3D [53], due to the higher dependence on temporal convolutions.

### 4.3 Action recognition performance and comparisons

Table 4 compares performances of different methods in the literature, across the two datasets for action recognition - two protocols for the NTU RGB+D and the NWUCLA. The standard performance measure used for this task and datasets is classification accuracy, estimated according to the protocols (training and testing splits) reported in the respective works we compare to. The first part of the table (indicated by  $\times$  symbol) refers to unsupervised methods, which achieve surprisingly high results even without relying on labels in learning representations.

The second part refers to supervised methods (indicated by  $\triangle$ ), divided according to the modalities used for training and testing. Here, we report the performance of the separate RGB and depth (with and without bottleneck) streams trained in step 1 (rows #7 and #8). The small increase in performance is probably due to the extra training steps with small learning rate, after initialized with the bottleneck version trained on the mini-NTU (used for the ablation study). Importantly, the depth stream with bottleneck represents the teacher network used for the hallucination learning. We expect our final model to perform better than the one trained on RGB only, whose accuracy constitutes a lower bound for the usefulness of our hallucination model. The values reported for our step 1 models for the NW-UCLA dataset, *i.e.* the RGB and depth streams, refer to the fine-tuning of our NTU model. In contrast with [11], and for clearer analysis, the TwoStream setup is always not finetuned. Its accuracy represents an upper bound for the final model, which will not rely on depth data at test time. We have experimented training using pre-trained ImageNet weights instead of the NTU, but it led to lower accuracy.

The last part of the table (indicated by  $\square$ ) reports the performance of methods in the privileged information framework, thus directly comparable to ours. The performance

values that refer to the Hoffman *et al.* method [10] (row #20 of Table 4) are taken from the implementation and experiments in [11]. Step 3 and 4 of [11] (row #21 and #22) refer to the two-stream model after the hallucination learning, and its fine-tuning, respectively. We note that, for simplicity, the results of ADMD Two-Stream models are merely the outcome of the average of the two streams’ logits, and are not subject to any fine-tuning, which means that they are directly comparable with row #21. In addition, results of row #23 correspond to the hallucination stream only.

We note that the hallucination stream (row #23) manage to recover and surpass the depth teacher stream (row #8) for the NWUCLA dataset (83.94% compared to 71.09%), while for the NTU p1 (67.57%) and p2 (71.80%) protocols is around 4% below the respective teacher (71.87% and 75.32%). Nevertheless, when combined with the RGB stream, it performs better (NTU p2 - 81.50%) or comparable (NTU p1 - 73.11%) to the fine-tuned model presented in [11]. Since the RGB stream is performing equally well in this work and in [11], we can conclude that the gains in performance are due to better hallucination features.

### 4.4 Object recognition performance and comparisons

Table 5 illustrates the main results obtained for NYUD dataset for the object recognition task.

As opposed to action recognition, depth information is often noisy here (cfr. Fig. 5 - chair and lamp), probably due to the small resolution of the bounding box crops. Depth alone is in fact performing worse than RGB alone (more than 10% gap). Still, the amount of *complementary information* carried by the two modalities is able, when fused in the TwoStream model, to boost recognition accuracy by more than 5 percentage points, despite the poor depth performance (RGB  $\rightarrow$  52.90%, Depth  $\rightarrow$  40.19%  $\Rightarrow$  TwoStream  $\rightarrow$  57.39%).

It is well established that ensemble methods tend to outperform their single-model counterparts: an ensemble of two CNNs, each trained started from a different initialization, outperforms either model independently [60]. Since, in principle, the proposed ADMD strategy is the combination of an RGB model trained using a standard supervised approach and *another* adversarially trained RGB model, we additionally compare our approach to an ensemble of RGB classifiers (third line of Table 5). Interestingly, despite starting from a two relatively high single-stream performances, the fusion process of two RGB networks only marginally increases the final accuracy (RGB1  $\rightarrow$  53.19%, RGB2  $\rightarrow$  52.60%  $\Rightarrow$  Ensemble  $\rightarrow$  54.14%).

As for the task of action recognition, we found that fine-tuning the fused streams does not always bring significant improvements, as opposed to [11], were the architecture features cross-stream multiplier connections, which need to be trained in an additional step. Fine-tuning with the strategy proposed by Neverova *et al.* [19] looks slightly more effective, since ModDrop introduces a light dropout at the input layers, both on the images and on the whole modalities. The resulting model is tested in both the original setup proposed in [19], namely by blanking out the depth stream, and by simply using RGB predictions. The

TABLE 4

Classification accuracies and comparisons with the state of the art for video action recognition. Performances referred to the several steps of our approach (ours) are highlighted in bold.  $\times$  refers to comparisons with unsupervised learning methods.  $\triangle$  refers to supervised methods: here train and test modalities coincide.  $\square$  refers to privileged information methods: here training exploits RGB+D data, while test relies on RGB data only. The 3rd column refers to cross-subject and the 4th to the cross-view evaluation protocols on the NTU dataset. The results reported on the other two datasets are for the cross-view protocol.

#	Method	Test Mods.	NTU (p1)	NTU (p2)	NW-UCLA
1	Luo [54]	Depth	66.2%	-	-
2	Luo [54]	RGB	56.0%	-	-
3	Rahmani [55]	RGB	-	-	78.1%
4	HOG-2 [56]	Depth	32.4%	22.3%	-
5	Action Tube [57]	RGB	-	-	61.5%
6	Depth stream [11]	Depth	70.44%	75.16%	72.38%
7	ADMD - Depth stream	Depth	70.53%	76.47%	-
8	ADMD - Depth stream w/ bott.	Depth	71.87%	75.32%	71.09%
9	[11] - RGB stream	RGB	66.52%	80.01%	85.22%
10	ADMD - RGB stream	RGB	67.95%	80.01%	85.87%
11	Deep RNN [15]	Joints	56.3%	64.1%	-
12	Deep LSTM [15]	Joints	60.7%	67.3%	-
13	Sharoudy [15]	Joints	62.93%	70.27%	-
14	Kim [58]	Joints	74.3%	83.1%	-
15	Sharoudy [5]	RGB+D	74.86%	-	-
16	Liu [6]	RGB+D	77.5%	84.5%	-
17	Rahmani [59]	Depth+Joints	75.2	83.1	-
18	Two-stream, step 2 [11]	RGB+D	79.73%	81.43%	88.87%
19	ADMD - Two-stream (no finetune)	<b>RGB+D</b>	<b>77.74%</b>	<b>85.49%</b>	<b>89.93%</b>
20	Hoffman <i>et al.</i> [10]	RGB	64.64%	-	83.30%
21	Hallucination model, step 3 [11]	RGB	71.93%	74.10%	76.30%
22	Hallucination model, step 4 [11]	RGB	73.42%	77.21%	86.72%
23	ADMD - Hall. stream alone	<b>RGB</b>	<b>67.57%</b>	<b>71.80%</b>	<b>83.94%</b>
24	ADMD - Hall. two-stream model	<b>RGB</b>	<b>73.11%</b>	<b>81.50%</b>	<b>91.64%</b>

TABLE 5  
Object Recognition

Method	Trained on	Tested on	Accuracy
Depth alone	Depth	Depth	40.19%
RGB alone	RGB	RGB	52.90%
RGB ensemble	RGB	RGB	54.14%
TwoStream (average logits)	RGB+D	RGB+D	57.39%
TwoStream after finetuning	RGB+D	RGB+D	58.73%
ModDrop [19] (finetuned from TwoStream)	RGB+D	RGB+D	58.93%
ModDrop [19]	RGB+D	RGB+blankD	47.86%
ModDrop [19]	RGB+D	RGB	53.73%
Autoencoder	RGB+D	RGB	50.52%
Ours (naive adversarial)	RGB+D	RGB	50.81%
<b>Ours (ADMD)</b>	RGB+D	RGB	<b>57.52%</b>

latter scheme slightly improves the performance of the RGB stream, possibly thanks to dropout. However, although the model shows more robustness to missing depth at test time, it clearly fails to extract any monocular depth cue.

Another interesting comparison we perform is the following: we train a cross-modal autoencoder with an L2 loss, in order to reconstruct depth maps from RGB. The encoder-decoder architecture consists in the very same RGB ResNet-50 for the encoder, and in 5 stacked deconvolutional blocks intertwined with batch-norm layers for the decoder. At test time, when depth is not available, we provide RGB frames to the autoencoder, which reconstructs the missing modality to feed the corresponding branch of the TwoStream architecture. The performance of this setup is quite poor. We

observe that the autoencoder easily overfits the training set, generating high quality depth maps for the training set, while it performs very poorly for the test set.





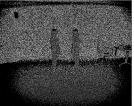


Eventually, we tested our adversarial scheme in two different setups: i) the naive setup where the discriminator  $D$  is assigned the binary task only, and ii) the ADMD setup, where the discriminator is also assigned the classification task. While the former performs as the autoencoder, the latter is able to fully recover the accuracy of the TwoStream model, being only slightly below that of the fine-tuned model.

#### 4.5 Inference with noisy depth



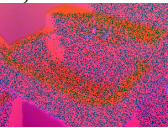
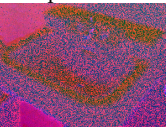



In real test scenarios, it is often the case that we can only access *noisy* depth data. In this section, we address two questions: i) how much such noisy data can degrade the performance of a multimodal setup? ii) At which level of noise does it become favorable to hallucinate the depth modality with respect to using the teacher model (TwoStream) with noisy depth data?

The depth sensor used in the NTU dataset (Kinect), is an IR emitter coupled with an IR camera, and has very complex noise characterization comprising at least 6 different sources [61]. It is beyond the scope of this work to investigate noise models affecting the depth channel, so, for our analysis we choose the most influencing one, i.e., multiplicative speckle noise. Hence, we inject Gaussian noise in the depth images  $I$  in order to simulate speckle noise:  $I = I * n$ ,  $n \sim \mathcal{N}(1, \sigma)$ . Table 6 shows how performances of our TwoStream network degrade when depth is corrupted with such Gaussian noise

TABLE 6  
Accuracy values for the two-stream model trained on RGB and depth, and tested with RGB and noisy depth data.

NTU RGB+D action dataset - ADMD performance is 81.50%.							
							
$\sigma^2$	no noise	$10^{-3}$	$10^{-2}$	$10^{-1}$	$10^0$	$10^1$	void
TwoStream	85.49%	85.52%	82.05%	68.99%	2.16%	3.35%	8.55%

NYUD object dataset - ADMD performance is 57.52%.							
							
$\sigma^2$	no noise	$10^{-3}$	$10^{-2}$	$10^{-1}$	$10^0$	$10^1$	void
TwoStream	58.73%	58.68%	58.23%	57.18%	48.27%	28.40%	47.44%
ModDrop [19]	58.93%	58.89%	58.56%	57.49%	48.90%	25.95%	47.86%

with increasing variance (NTU cross-view protocol and NYUD). Results show that accuracy significantly decreases with respect to the one guaranteed by our hallucination model (81.50% - row #24) in Table 4, even with low noise variance of  $\sigma^2=10^{-1}$ . For the task of object recognition, we can see that ModDrop [19] is slightly more resilient to depth corruption than the simple TwoStream, since fine-tuned with noise (dropout) in the input layer.

This experiment shows, in conclusion, that *ADMD can not only deal with a missing modality, but also with a noisy one*. In an online scenario, the discriminator  $D$ , trained in step 2, can give an indication on when to operatively switch from TwoStream to ADMD, that is, when to substitute the depth branch with the hallucination. When training reaches equilibrium,  $D$  is maximally fooled by the features generated by  $H$ , and cannot distinguish them from those encoded by  $E_d$ . In practice, this means that the predicted probability for the fake class (last class in  $\hat{y}$ , eq. 1) is  $p(\hat{y} = C + 1) \approx .5$  on




Fig. 6. Discriminator confidence at predicting ‘fake’ label as a function of noise in the depth frames. The more corrupted the frame, the more confident  $D$ , and the lower the accuracy of the TwoStream model (NYUD dataset).

average. However, when features computed from corrupted depth start to flow inside  $D$ , its prediction for the fake class starts to be more and more confident. Figure 6 plots the behavior of  $D$  as noise increases, together with accuracy of the TwoStream model. There is a clear turning point in both accuracy and confidence, which can be employed in practice to decide when to switch from  $E_d$  to  $H$  *i.e.* when to drop depth as a modality and start using monocular depth features extracted from RGB.

#### 4.6 Discussion

Some interesting points arise from the analysis of our findings, which we summarize in the following.

**1. RGB and depth actually carry complementary information.** As a matter of fact, the TwoStream setup always provides a surprisingly better accuracy than the two streams alone. As additional evidence, a multimodal ensemble (*i.e.* the TwoStream) performs better than a mono-modal ensemble (Table 5), despite the lower accuracy of one of its single-stream components (either depth or RGB, depending on task and dataset).

**2. There is (monocular) depth information in RGB images.** This is evident from the fact that the hallucination stream often recovers and sometimes surpasses the accuracy of its depth-based teacher network. Besides, fusing hallucination and RGB streams always bring the benefits, as fusing RGB and Depth.

**3. Standard supervised learning has limitations in extracting information.** In fact, given the evidence that there is depth information to exploit in RGB images, minimizing cross-entropy loss is not enough to fully extract it. For that we need a student-teacher adversarial framework. This has an interesting parallel in *adversarial network compression* [36], where the performance of a fully supervised small network can be boosted by adversarial training against a high-capacity (and better performing) teacher net. In [36], it is also observed that the student can surpass the teacher in some occasions.

**4. Adversarial training alone only is not enough.** The naive discriminator trained for the binary task

(real/generated) is not sufficient to force the hallucination network to produce discriminative features. The auxiliary discriminative task is necessary to extract monocular depth cues which are also discriminative for a given task (on the other hand, the auxiliary task only is not enough, as suggested by the performance of the RGB ensemble).

## 5 CONCLUSIONS

In this work, we have introduced a novel technique to exploit additional information, in the form of depth images at training time, to improve RGB only models at test time. This is done by adversarially training a hallucination network which learns from a teacher depth stream how to encode monocular depth features from RGB frames. The proposed approach outperforms previous ones in the privileged information scenario in the tasks of object classification and action recognition on three different datasets. Additionally, the hallucination framework is shown to be very effective in cases where depth is noisy. Code will be made available at <https://github.com/pmorerio/admd>

## ACKNOWLEDGMENTS

The authors would like to thank Riccardo Volpi for useful discussion on adversarial training and GANs.

## REFERENCES


- [1] E. J. Gibson and R. D. Walk, "The "visual cliff","" *Scientific American*, vol. 202, no. 4, pp. 64–71, 1960. [Online]. Available: <http://www.jstor.org/stable/24940447>
- [2] M. R. Watson and J. T. Enns, "Depth perception," *Encyclopedia of Human Behavior*, 2012.
- [3] P. Servos, "Distance estimation in the visual and visuomotor systems," *Experimental Brain Research*, vol. 130, no. 1, pp. 35–47, Jan 2000. [Online]. Available: <https://doi.org/10.1007/s002210050004>
- [4] M. Firman, "Rgb-d datasets: Past, present and future," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2016, pp. 19–31.
- [5] A. Shahroudy, T.-T. Ng, Y. Gong, and G. Wang, "Deep multimodal feature analysis for action recognition in rgb+ d videos," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.
- [6] J. Liu, N. Akhtar, and A. Mian, "Viewpoint invariant action recognition using rgb-d videos," *arXiv preprint arXiv:1709.05087*, 2017.
- [7] S. Gupta, R. Girshick, P. Arbeláez, and J. Malik, "Learning rich features from rgb-d images for object detection and segmentation," in *European Conference on Computer Vision*, 2014, pp. 345–360.
- [8] C. Hazirbas, L. Ma, C. Domokos, and D. Cremers, "Fusenet: Incorporating depth into semantic segmentation via fusion-based cnn architecture," in *Asian Conference on Computer Vision*. Springer, 2016, pp. 213–228.
- [9] V. Vapnik and A. Vashist, "A new learning paradigm: Learning using privileged information," *Neural networks*, vol. 22, no. 5, pp. 544–557, 2009.
- [10] J. Hoffman, S. Gupta, and T. Darrell, "Learning with side information through modality hallucination," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 826–834.
- [11] N. Garcia, P. Morerio, and V. Murino, "Modality distillation with multiple stream networks for action recognition," *European Conference on Computer Vision*, 2018.
- [12] D. Lopez-Paz, L. Bottou, B. Schölkopf, and V. Vapnik, "Unifying distillation and privileged information," *Proceedings of the International Conference on Learning Representations (ICLR)*, 2016.
- [13] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in neural information processing systems*, 2014, pp. 2672–2680.
- [14] S. Arora, A. Risteski, and Y. Zhang, "Do GANs learn the distribution? some theory and empirics," in *International Conference on Learning Representations*, 2018. [Online]. Available: <https://openreview.net/forum?id=BjehNfW0>
- [15] A. Shahroudy, J. Liu, T.-T. Ng, and G. Wang, "Ntu rgb+ d: A large scale dataset for 3d human activity analysis," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 1010–1019.
- [16] J. Wang, X. Nie, Y. Xia, Y. Wu, and S.-C. Zhu, "Cross-view action modeling, learning and recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 2649–2656.
- [17] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *Deep Learning and Representation Learning Workshop: NIPS 2014*, 2014.
- [18] C. Feichtenhofer, A. Pinz, and R. P. Wildes, "Spatiotemporal multiplier networks for video action recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 4768–4777.
- [19] N. Neverova, C. Wolf, G. W. Taylor, and F. Nebout, "Moddrop: Adaptive multi-modal gesture recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 38, no. 8, pp. 1692–1706, 2016. [Online]. Available: <https://doi.org/10.1109/TPAMI.2015.2461544>
- [20] Z. Luo, L. Jiang, J.-T. Hsieh, J. C. Niebles, and L. Fei-Fei, "Graph distillation for action detection with privileged information," in *European Conference on Computer Vision*, 2018.
- [21] Y. Zhang, T. Xiang, T. M. Hospedales, and H. Lu, "Deep mutual learning,"
- [22] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," in *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, vol. 1. IEEE, 2005, pp. 886–893.
- [23] H. Wang and C. Schmid, "Action recognition with improved trajectories," in *Proceedings of the IEEE international conference on computer vision*, 2013, pp. 3551–3558.
- [24] I. Laptev, M. Marszalek, C. Schmid, and B. Rozenfeld, "Learning realistic human actions from movies," in *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*. IEEE, 2008, pp. 1–8.
- [25] E. S. Ye and J. Malik, "Object detection in rgb-d indoor scenes," *Technical Report of University of California at Berkeley*, 2013.
- [26] S. Tang, X. Wang, X. Lv, T. X. Han, J. Keller, Z. He, M. Skubic, and S. Lao, "Histogram of oriented normal vectors for object recognition with a depth sensor," in *Asian conference on computer vision*. Springer, 2012, pp. 525–538.
- [27] A. Janoch, S. Karayev, Y. Jia, J. T. Barron, M. Fritz, K. Saenko, and T. Darrell, "A category-level 3d object dataset: Putting the kinect to work," in *Consumer depth cameras for computer vision*. Springer, 2013, pp. 141–165.
- [28] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei, "Large-scale video classification with convolutional neural networks," in *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2014, pp. 1725–1732.
- [29] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri, "Learning spatiotemporal features with 3d convolutional networks," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 4489–4497.
- [30] X. Wang, R. Girshick, A. Gupta, and K. He, "Non-local neural networks," *arXiv preprint arXiv:1711.07971*, 2017.
- [31] A. Wang, J. Lu, J. Cai, T.-J. Cham, and G. Wang, "Large-margin multi-modal deep learning for rgb-d object recognition," *IEEE Transactions on Multimedia*, vol. 17, no. 11, pp. 1887–1898, 2015.
- [32] K. Simonyan and A. Zisserman, "Two-stream convolutional networks for action recognition in videos," in *Advances in neural information processing systems*, 2014, pp. 568–576.
- [33] D. Tran, H. Wang, L. Torresani, J. Ray, Y. LeCun, and M. Paluri, "A closer look at spatiotemporal convolutions for action recognition," in *Proceedings of the IEEE International Conference on Computer Vision*, 2018.
- [34] X. Wang, R. Girshick, A. Gupta, and K. He, "Non-local neural networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [35] X. Chen, Y. Duan, R. Houthoofd, J. Schulman, I. Sutskever, and P. Abbeel, "Infogan: Interpretable representation learning by information maximizing generative adversarial nets," in *Advances in neural information processing systems*, 2016, pp. 2172–2180.

- [36] V. Belagiannis, A. Farshad, and F. Galasso, "Adversarial network compression," *arXiv preprint arXiv:1803.10750*, 2018.
- [37] R. Volpi, P. Morerio, S. Savarese, and V. Murino, "Adversarial feature augmentation for unsupervised domain adaptation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [38] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [39] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein generative adversarial networks," in *International Conference on Machine Learning*, 2017, pp. 214–223.
- [40] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen, "Improved techniques for training gans," in *Advances in Neural Information Processing Systems*, 2016, pp. 2234–2242.
- [41] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," *arXiv preprint arXiv:1511.06434*, 2015.
- [42] M. Mirza and S. Osindero, "Conditional generative adversarial nets," *arXiv preprint arXiv:1411.1784*, 2014.
- [43] S. Roheda, B. S. Riggan, H. Krim, and L. Dai, "Cross-modality distillation: A case for conditional generative adversarial networks," *arXiv preprint arXiv:1807.07682*, 2018.
- [44] "How to train a gan? tips and tricks to make gans work," <https://github.com/soumith/ganhacks>.
- [45] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [46] L. Sun, K. Jia, D.-Y. Yeung, and B. E. Shi, "Human action recognition using factorized spatio-temporal convolutional networks," in *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 4597–4605.
- [47] A. Eitel, J. T. Springenberg, L. Spinello, M. Riedmiller, and W. Burgard, "Multimodal deep learning for robust rgb-d object recognition," in *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*. IEEE, 2015, pp. 681–687.
- [48] S. Gupta, R. Girshick, P. A. Aez, and J. Malik, "Learning rich features from rgb-d images for object detection and segmentation," in *European Conference on Computer Vision (ECCV)*, 2014.
- [49] N. Silberman, D. Hoiem, P. Kohli, , and R. Fergus, "Indoor segmentation and support inference from rgb-d images," in *European Conference on Computer Vision (ECCV)*, 2012.
- [50] E. Tzeng, J. Hoffman, K. Saenko, and T. Darrell, "Adversarial discriminative domain adaptation," in *CVPR*, 2017.
- [51] P. Morerio, J. Cavazza, and V. Murino, "Minimal-entropy correlation alignment for unsupervised deep domain adaptation," in *International Conference on Learning Representations*, 2018. [Online]. Available: <https://openreview.net/forum?id=rJWechg0Z>
- [52] K. He, X. Zhang, S. Ren, and J. Sun, "Delving deep into rectifiers: Surpassing human-level performance on imagenet classification," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 1026–1034.
- [53] J. Carreira and A. Zisserman, "Quo vadis, action recognition? a new model and the kinetics dataset," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [54] Z. Luo, B. Peng, D.-A. Huang, A. Alahi, and L. Fei-Fei, "Unsupervised learning of long-term motion dynamics for videos," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, no. EPFL-CONF-230240, 2017.
- [55] H. Rahmani, A. Mian, and M. Shah, "Learning a deep model for human action recognition from novel viewpoints," *IEEE transactions on pattern analysis and machine intelligence*, vol. 40, no. 3, pp. 667–681, 2018.
- [56] E. Ohn-Bar and M. M. Trivedi, "Joint angles similarities and hog2 for action recognition," in *Computer vision and pattern recognition workshops (CVPRW), 2013 IEEE conference on*. IEEE, 2013, pp. 465–470.
- [57] G. Gkioxari and J. Malik, "Finding action tubes," in *Computer Vision and Pattern Recognition (CVPR), 2015 IEEE Conference on*. IEEE, 2015, pp. 759–768.
- [58] T. Soo Kim and A. Reiter, "Interpretable 3d human action analysis with temporal convolutional networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2017, pp. 20–28.
- [59] H. Rahmani and M. Bennamoun, "Learning action recognition model from depth and skeleton videos," in *Proceedings of the IEEE*

*Conference on Computer Vision and Pattern Recognition*, 2017, pp. 5832–5841.

- [60] J. Guo and S. Gould, "Deep cnn ensemble with data augmentation for object detection," *CoRR*, vol. abs/1506.07224, 2015.

- [61] T. Mallick, P. P. Das, and A. K. Majumdar, "Characterizations of noise in kinect depth images: A review," *IEEE Sensors Journal*, vol. 14, no. 6, pp. 1731–1740, June 2014.



**Nuno C. Garcia** is a PhD fellow at the Pattern Analysis and Computer Vision department (PAVIS) in Istituto Italiano di Tecnologia (IIT) and at the Università degli studi di Genova. He was a data analytics consultant at Deloitte Portugal and a data engineer at Miniclip. He received his M. Sc. degree in Computer Engineering from the University of Beira Interior (Portugal) in 2015 (magna cum laude). His research interests are computer vision and machine learning.



**Pietro Morerio** Pietro Morerio received his M. Sc. in Theoretical Physics from the University of Milan (Italy) in 2010 (summa cum laude). He was Research Fellow at the University of Genoa (Italy) from 2011 to 2012, working in Video Analysis for Interactive Cognitive Environments. He pursued a PhD degree in Computational Intelligence at the same institution in 2016. Currently he is a Postdoctoral Researcher at Istituto Italiano di Tecnologia (IIT), Pattern Analysis and Computer Vision (PAVIS) department, his research including machine learning, deep learning and computer vision,

with a particular focus on multimodal learning.



**Vittorio Murino** received the Laurea degree in electronic engineering and the Ph.D. degree in electronic engineering and computer science from the University of Genova, Italy, in 1989 and 1993, respectively. He is a full professor at the University of Verona, Italy since 2000. Since 1998, he has been with the University of Verona where he held the Chair of the Department of Computer Science of this University from 2001 to 2007. He is currently with the Istituto Italiano di Tecnologia, as director of Pattern Analysis and

Computer Vision (PAVIS) department, involved in computer vision, machine learning, and image analysis activities. He has co-authored over 400 papers published in refereed journals and international conferences. His current research interests include computer vision, pattern recognition, and machine learning, more specifically, statistical and probabilistic techniques for image and video processing, with applications on (human) behavior analysis and related applications such as video surveillance, biomedical imaging, and bioinformatics. He is also an Associate Editor of *Computer Vision and Image Understanding*, *Machine Vision & Applications*, and *Pattern Analysis and Applications* journals. He is also an IAPR Fellow since 2006.