

# Guided Weak Supervision for Action Recognition with Scarce Data to Assess Skills of Children with Autism

Prashant Pandey,<sup>1</sup> Prathosh AP,<sup>2</sup> Manu Kohli,<sup>3</sup> Josh Pritchard<sup>4</sup>

<sup>1</sup>IIT Delhi, <sup>2</sup>IIT Delhi

<sup>3</sup>IIT Delhi, <sup>4</sup>Florida Institute of Technology

getprashant57@gmail.com, prathoshap@iitd.ac.in, manu.kohli@cogniable.tech, josh@factari.com

## Abstract

Diagnostic and intervention methodologies for skill assessment of autism typically requires a clinician repetitively initiating several stimuli and recording the child’s response. In this paper, we propose to automate the response measurement through video recording of the scene following the use of Deep Neural models for human action recognition from videos. However, supervised learning of neural networks demand large amounts of annotated data that is hard to come by. This issue is addressed by leveraging the ‘similarities’ between the action categories in publicly available large-scale video action (source) datasets and the dataset of interest. A technique called Guided Weak Supervision is proposed, where every class in the target data is matched to a class in the source data using the principle of posterior likelihood maximization. Subsequently, classifier on the target data is re-trained by augmenting samples from the matched source classes, along with a new loss encouraging inter-class separability. The proposed method is evaluated on two skill assessment autism datasets, SSBD (Sundar Rajagopalan, Dhall, and Goecke 2013) and a real world Autism dataset comprising 37 children of different ages and ethnicity who are diagnosed with autism. Our proposed method is found to improve the performance of the state-of-the-art multi-class human action recognition models in-spite of supervision with scarce data.

## Introduction

Autism is a complex neuro-developmental disorder that manifests in children during preschool years (Rapin 1991; CDC 2019) as deficits in communication, social skills and stereotypical repetitive behavior. In the last two decades, the prevalence rate of Autism Spectrum Disorder (ASD) has grown by more than 150% (CDC 2019). It is well-established that early intervention services modelled on behavior therapies yield the best outcomes for children diagnosed with autism (Estes et al. 2015) leading to significant societal benefits and cost savings. However, resource and expert scarcity in low resource settings result in delay in the initiation of the treatment process due to lack of identification of the disorder. Families have limited access to diagnostic and evidence-based treatment options because of affordability, lack of insurance support and non-availability of physical infrastructure. In the developing countries, there

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

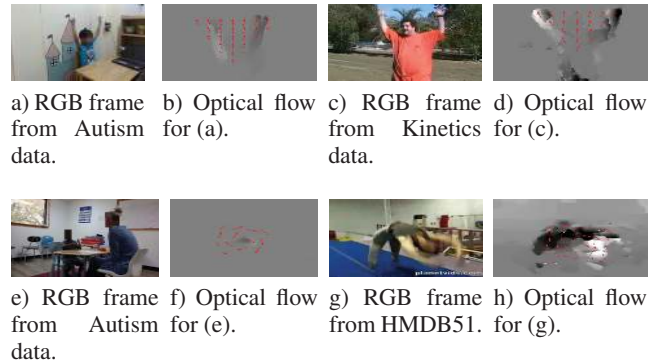


Figure 1: Illustration of the similarity between the action classes in two different datasets in the optical flow domain. One representative RGB and Flow frame is depicted in each case: The directional closeness of the optical flow frames can be observed despite being from unrelated classes. (a)-(b) ‘Arms up’ action in Autism data is close to (c)-(d) ‘Jumping jacks’ action in Kinetics (Kay et al. 2017) data. (e)-(f) ‘Rolly polly’ action in Autism data is close to (g)-(h) ‘Flic flac’ action in HMDB51 (Kuehne et al. 2011) data.

are entire swaths of children in dire need of treatment who remain untreated and unseen. Much of this strain and inaccessibility could be alleviated by incorporating technology to assist in early screening and automate the assessment and initial treatment planning process, which is a prerequisite to deliver individualized behavior treatment to children. Skill assessment processes for autism typically involves invoking instructions to a child, monitoring and recording their responses as they occur. This requires a trained clinician to engage with the child, perform mundane repetitive tasks such as recording the child’s observation and human action responses to fixed set of stimuli. With the advent of tremendously powerful modern Deep learning techniques, one can hope to automate a lot of such tasks bringing affordability and improved access in value chain of autism screening, diagnosis and behavioral treatment activities while reducing the dependence on trained clinicians. Specifically, in this paper, we examine the application of human action recognition from video recordings for tracking the physical behavior of children diagnosed with ASD or otherwise to build cognitive

and functional assessments.

## Motivation

Human action recognition is typically solved in a supervised machine learning setting. Most of the successful models employ Convolutional Neural Networks (CNN) as their backbone. Two stream networks (Simonyan and Zisserman 2014), 3-Dimensional (3D) Convolutional Neural Networks (Tran et al. 2015) and Convolutional long short-term memory (LSTM) networks (Donahue et al. 2015) are some of the state-of-the-art action recognition approaches. Inflated 3D CNN (Carreira and Zisserman 2017) or I3D is found to be one of the top performers on the standard benchmarks like UCF101 (Soomro, Zamir, and Shah 2012) and HMDB51 (Kuehne et al. 2011) datasets. Temporal Segment Networks (Wang et al. 2016) or TSN is another example of two stream networks which has better ability to model long range temporal dependencies.

The objective of this work is to apply human action recognition algorithms to evaluate the responses of the children with autism, to measure stimulus in the area of imitation, listener response and motor skills. To accomplish this using already established methods require large amounts of expert annotated data corresponding to the particular classes of actions to be recognized. The process of data collection and annotation is non-trivial because of non-availability of expert annotators, lack of co-operation from children and is also very time consuming. Despite these limitations, there is abundant availability of large-scale public datasets that contain thousands of annotated video clips corresponding to hundreds of action classes. Further, human action classes share a lot of similarity (in a well-defined feature space like optical flow (Horn and Schunck 1981)) in spite of being disjoint as shown in Figure 1. For example, intuitively, the action classes ‘playing piano’ and ‘typing the keyboard’ can be thought similar in a suitable feature space. Motivated by these observations, in this paper we address the following question - *Given a target data distribution with few annotated samples, can the availability of a large-scale annotated source dataset with ‘similar’ attributes as the target data, be leveraged, to improve the generalization abilities of classifiers trained on target data?* Specifically, we attempt to weakly supervise the task of human action recognition in a guided way using large-scale publicly available datasets.

## Related Work

State-of-the-art action recognition models (Zolfaghari, Singh, and Brox 2018; Lin, Gan, and Han 2019; Ghadiyaram, Tran, and Mahajan 2019; Carreira and Zisserman 2017; Wang et al. 2016) are deep neural networks which overfits easily to the dataset with fewer samples leading to poor generalization. Few-Shot Learning (FSL) action recognition methods (Yang, He, and Porikli 2018; Zhu and Yang 2018; Wang, Zhou, and Qiao 2018; Mandal et al. 2019) are characterised by different labels between source and target but a similar feature space. FSL evaluates on novel classes with limited training examples but these novel classes are sampled from the same dataset. The large scale public

datasets are not good candidate for source in FSL as there is significant domain shift that exists between the source and target data.

## Proposed Method

For a given video, there exists transformations such as optical flow, that are non-unique mappings of the video space. This suggests that given multiple disjoint set of action classes, there can be spaces (such as flow) where a given pair of action classes may lie ‘close’ albeit they represent different semantics in the RGB-space. For example, the optical flow characteristics of a ‘baseball-strike’ class and ‘cricket-slog’ class can be imagined to be close. Further, there exists large-scale public datasets (e.g., Kinetics (Kay et al. 2017)) that encompasses a large number of annotated videos for several action classes. Thus, if one can find the classes in the public datasets that are ‘close’ to a given class in the dataset of interest, then the videos from the public dataset can be potentially used for augmentation resulting in regularization. In the subsequent sub-sections, we will formalize the aforementioned idea and describe a procedure to find the closer classes and use it for data augmentation. Let  $\mathcal{X}$  denote the sample space encompassing the elements of transformed videos (e.g., optical flow). Let  $\mathbf{P}^s(\mathbf{x}_s)$  and  $\mathbf{P}^t(\mathbf{x}_t)$  be two distributions on  $\mathcal{X}$  called the source and target distributions respectively. Suppose a semantic labeling scheme is defined both on  $\mathbf{P}^s(\mathbf{x}_s)$  and  $\mathbf{P}^t(\mathbf{x}_t)$ . That is, let  $\mathcal{Y}_s = \{\mathbf{y}_s^1, \mathbf{y}_s^2, \dots, \mathbf{y}_s^N\}$  and  $\mathcal{Y}_t = \{\mathbf{y}_t^1, \mathbf{y}_t^2, \dots, \mathbf{y}_t^M\}$  be the source and target class labels that are assigned for the samples of  $\mathbf{P}^s(\mathbf{x}_s)$  and  $\mathbf{P}^t(\mathbf{x}_t)$  respectively which in-turn defines the joint distributions  $\mathbf{P}^s(\mathbf{x}_s, \mathbf{y}_s)$  and  $\mathbf{P}^t(\mathbf{x}_t, \mathbf{y}_t)$ .  $N$  and  $M$  are the respective number of source and target classes. Let  $\mathcal{D}_s = \{(\mathbf{x}_s, \mathbf{y}_s)\}$  and  $\mathcal{D}_t = \{(\mathbf{x}_t, \mathbf{y}_t)\}$  denote the tuples of samples drawn from the two joint distributions  $\mathbf{P}^s$  and  $\mathbf{P}^t$ , respectively. Suppose a parametric discriminative classifier (Deep Neural Network) is learned using  $\mathcal{D}_t$  to obtain estimate of the conditional distribution  $\mathbf{P}_\theta^t(\mathbf{y}_t|\mathbf{x}_t)$  where  $\theta$  denotes the parameters of the neural network. With these notations, we consider the case where the cardinality of  $\mathcal{D}_t$  is much less than that of  $\mathcal{D}_s$  implying that the amount of supervised data in the case of target distribution is much less than that of the source distribution. In such a case,  $\mathbf{P}_\theta^t(\mathbf{y}_t|\mathbf{x}_t)$  trained on  $\mathcal{D}_t$  is deemed to overfit and hence doesn’t generalize well. If there exists a  $\mathbf{y}_s^p \in \mathcal{Y}_s$  that is ‘close’ to  $\mathbf{y}_t^q \in \mathcal{Y}_t$ , then samples drawn from  $\mathbf{P}^s(\mathbf{x}_s|\mathbf{y}_s = \mathbf{y}_s^p)$  can be used to augment the class  $\mathbf{y}_t^q$  for re-training the model  $\mathbf{P}_\theta^t(\mathbf{y}_t|\mathbf{x}_t)$ . In the subsequent sub-section, we describe a procedure to find the ‘closest’  $\mathbf{y}_s^p \in \mathcal{Y}_s$ , given  $\mathbf{y}_t^q \in \mathcal{Y}_t$  and a model  $\mathbf{P}_\theta^t(\mathbf{y}_t|\mathbf{x}_t)$  trained on  $\mathcal{D}_t$ .

## Guided Weak Supervision

Videos lie in a very high dimensional space and are of variable length in general. Thus, standard vector distance metrics are not feasible to measure the closeness of two video objects. Further, the objective here is to quantify the distance between the classes as perceived by the discriminative model (classifier)  $\mathbf{P}_\theta^t(\mathbf{y}_t|\mathbf{x}_t)$  so that the data augmentation

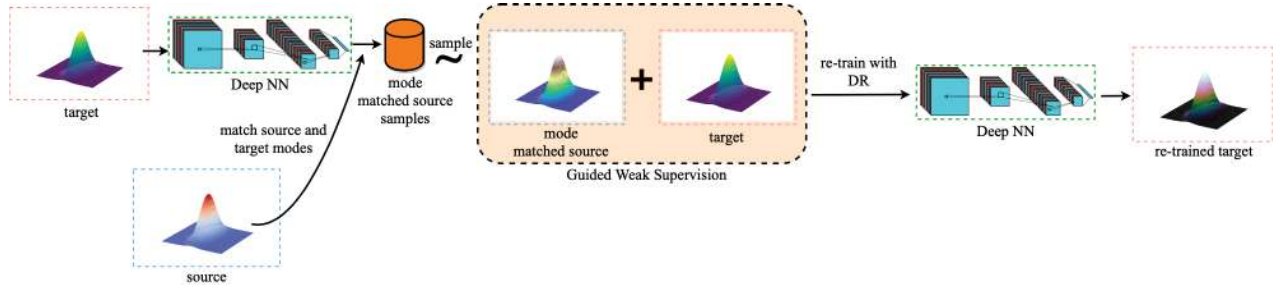


Figure 2: The overall procedure of the proposed method. A classifier trained on the target data is used to match the modes (classes) of the source data. Target classes are augmented with samples from the matched source classes (Guided Weak Supervision). Finally, the classifier is re-trained with the augmented target classes along with Directional Regularization loss.

is sensible. Thus, we propose to use the maximum posterior likelihood principle to define the closeness between two classes. Let  $\mathcal{X}_{(y_s=y_s^p)} = \{\mathbf{x}_{s1}, \mathbf{x}_{s2}, \dots, \mathbf{x}_{s1}\}$  denote the samples drawn from  $\mathbf{P}^s(\mathbf{x}_s | y_s = y_s^p)$ . Now  $\mathbf{P}_\theta^t(y_{tj} | \mathbf{x} = \mathbf{x}_{sj})$ ,  $j \in \{1, 2, \dots, 1\}$  denotes the posterior distribution of the target classes  $\mathcal{Y}_t$  given the  $j^{\text{th}}$  feature vector from the source class  $\mathcal{X}_{(y_s=y_s^p)}$ . With this, a joint posterior likelihood  $\mathcal{L}_{y_t|\mathbf{x}_s}$  of a class  $y_s^p$  can be defined as observing the target classes given a set of features  $\mathcal{X}_{(y_s=y_s^p)}$  drawn from a particular source class  $y_s^p$ . Mathematically,

$$\mathcal{L}_{y_t|\mathbf{x}_s} = \mathbf{P}_\theta^t(y_{t1}, y_{t2}, \dots, y_{t1} | \mathbf{x}_{s1}, \mathbf{x}_{s2}, \dots, \mathbf{x}_{s1}) \quad (1)$$

where  $\mathbf{x}_{sj}$ ,  $j \in \{1, 2, \dots, 1\}$ , are from the class  $y_s^p$ . If it is assumed that  $\mathbf{x}_{sj}$  are drawn i.i.d., one can express Eq. 1 as,

$$\mathcal{L}_{y_t|\mathbf{x}_s} = \prod_{j=1}^1 \mathbf{P}_\theta^t(y_{tj} | \mathbf{x}_{sj}) \quad (2)$$

The parameters  $\theta$  of the discriminative model created using  $\mathcal{D}_t$  are independent of  $\mathcal{X}_{(y_s=y_s^p)}$  and are fixed during the evaluation of  $\mathcal{L}_{y_t|\mathbf{x}_s}$ , which implies that  $y_{ti} | \mathbf{x}_{si}$  is independent of  $\mathbf{x}_{si} \forall i \neq j$  thus leading to Eq. 2. The posterior likelihood in Eq. 2 can be evaluated for every target class  $y_t = y_t^q$ ,  $q \in \{1, 2, \dots, M\}$ , denoted by  $\mathcal{L}_{y_t=y_t^q|\mathbf{x}_s}$  called the target-class posterior likelihood corresponding to the features from source class  $y_s^p$  under the learned target classifier  $\mathbf{P}_\theta^t(y_t | \mathbf{x}_t)$ . Mathematically,

$$\mathcal{L}_{y_t=y_t^q|\mathbf{x}_s} = \prod_{j=1}^1 \mathbf{P}_\theta^t(y_{tj} = y_t^q | \mathbf{x}_{sj}) \quad (3)$$

With this definition of the target-class posterior likelihood, we define the matched source class  $y_s^* | y_t^q \in \mathcal{Y}_s$  to a given target class  $y_t^q$  as follows:

$$y_s^* | y_t^q = \operatorname{argmax}_{y_s} \mathcal{L}_{y_t=y_t^q|\mathbf{x}_s} \quad (4)$$

Note that the definition of  $\mathcal{L}_{y_t=y_t^q|\mathbf{x}_s}$  is specific to a source-target class pair and therefore all  $\mathbf{x}_{sj}$  in the objective function of the optimization problem in Eq. 4 comes from a particular source class. Thus, one can employ the discriminative classifier trained on the target data to find out the ‘closest’

matching source class as the one that maximizes the posterior likelihood of observing that class as the given target class under the classifier. Every class in the joint distribution can be looked as a ‘mode’ and the goal here is to match the classes (‘modes’) in the joint distribution of the source and target distributions. Figure 1 demonstrates the idea of mode matching through examples. Optical flow frames of the target and the source classes have similar visual properties indicating the closeness. Once the matched source class is determined for every given target class, a set of source classes matched is defined as  $\mathcal{Y}_s^* = \{y_s^{*1}, y_s^{*2}, \dots, y_s^{*M}\}$ . Now, the discriminative classifier  $\mathbf{P}_\theta^t$  can be re-trained on the samples from the source dataset corresponding to  $\mathcal{Y}_s^*$  in a supervised way with class labels being the corresponding  $y_t^q$  for every  $y_s^*$ . This procedure thus increases the quantity and variety of the training data for  $\mathbf{P}_\theta^t$  and we call it as Guided Weak Supervision (GWS).

### Directional Regularization

The procedure of GWS described in the previous sub-section effectively changes the semantic meaning of the matched source classes to the semantic meaning of the target classes. Thus, it is possible to train a classifier on the source data to discriminate between the matched source classes  $\mathcal{Y}_s^* = \{y_s^{*1}, y_s^{*2}, \dots, y_s^{*M}\}$ . Suppose such a classifier is denoted by  $\mathbf{P}_\phi^s(y_s^* | \mathbf{x}_s)$ , where  $\phi$  are the model parameters. We assume that  $\mathbf{P}_\phi^s(y_s^* | \mathbf{x}_s)$  and  $\mathbf{P}_\theta^t(y_t | \mathbf{x}_t)$  have the same architectural properties. Also, it is assumed that the source dataset is larger in size and more diverse as compared to the target dataset. This implies that  $\mathbf{P}_\phi^s(y_s^* | \mathbf{x}_s)$  has better generalization abilities compared to  $\mathbf{P}_\theta^t(y_t | \mathbf{x}_t)$  (This fact is corroborated empirically in the experiment section). We propose to leverage this fact in improving the generalization capabilities of  $\mathbf{P}_\theta^t(y_t | \mathbf{x}_t)$  using  $\mathbf{P}_\phi^s(y_s^* | \mathbf{x}_s)$ . Further, during the training of  $\mathbf{P}_\phi^s(y_s^* | \mathbf{x}_s)$  with samples from  $\mathcal{Y}_s^*$ , it is desirable that the separation that is achieved between the classes in  $\mathcal{Y}_s^*$  under the classifier  $\mathbf{P}_\phi^s(y_s^* | \mathbf{x}_s)$  is ‘preserved’ during the training of  $\mathbf{P}_\theta^t(y_t | \mathbf{x}_t)$  with samples from  $\mathcal{Y}_s^*$ . We propose to accomplish the aforementioned properties by imposing a regularization term during the training of  $\mathbf{P}_\theta^t(y_t | \mathbf{x}_t)$ . Specifically, we propose to push the significant directions of the parameter matrix of  $\theta$  towards that of the parameter matrix



of  $\phi$ . Note that  $\phi$  is fixed during the training of  $\mathbf{P}_\theta^t(\mathbf{y}_t|\mathbf{x}_t)$ . Intuitively, this implies that the significant directions of the target parameters should follow that of the source parameters. Mathematically, let  $\mathbf{M}_\theta$  and  $\mathbf{M}_\phi$  be two square matrices formed by re-shaping (without any preference to particular dimensions) the parameters  $\theta$  and  $\phi$ , respectively. We perform an Eigenvalue Decomposition on  $\mathbf{M}_\theta$  and  $\mathbf{M}_\phi$  to obtain matrices of eigenvectors  $\mathbf{E}_\theta$  and  $\mathbf{E}_\phi$ , respectively. Let  $\hat{\mathbf{E}}$  denote the truncated versions of  $\mathbf{E}$  with first  $k$  significant (a model hyperparameter) eigenvectors. Under this setting, we desire the significant directions  $\hat{\mathbf{E}}_\theta$  and  $\hat{\mathbf{E}}_\phi$  to be aligned. Mathematically, if they are perfectly aligned, then

$$\hat{\mathbf{E}}_\theta^T \hat{\mathbf{E}}_\phi = \mathbf{I}_k \quad (5)$$

where  $\mathbf{I}_k$  is a  $k$ -dimensional identity matrix and  $\mathbf{T}$  denotes the transpose operation. Thus, any deviation from the condition laid in Eq. 5 is penalized by minimizing the Frobenius norm of the deviation. We term this as Directional Regularization (DR) denoted as  $L_{DR}$  given by the following equation:

$$L_{DR} = \|\hat{\mathbf{E}}_\theta^T \hat{\mathbf{E}}_\phi - \mathbf{I}_k\|_F \quad (6)$$

where  $\|\cdot\|_F$  denotes the Frobenius norm of a matrix. Note that this regularizer imposed on  $\theta$  during the training of  $\mathbf{P}_\theta^t(\mathbf{y}_t|\mathbf{x}_t)$  ensures that the directions of separating hyperplanes of the target classifier is encouraged to follow those of the source classifier trained with the matched classes. Thus the final objective function during the re-training of the target classifier is as follows:

$$L_{Total} = -\sum_{j=1}^1 \mathbf{y}_{tj} \log \hat{\mathbf{y}}_{tj} + \|\hat{\mathbf{E}}_\theta^T \hat{\mathbf{E}}_\phi - \mathbf{I}_k\|_F \quad (7)$$

where  $\hat{\mathbf{y}}_{tj}$  is the predicted target class. Thus in summary, given a small amount of data from the target distribution, the proposed method (1) trains a classifier on the target samples, (2) determines the closest classes from the source distribution to all the target classes using the target classifier, (3) trains a new (relatively robust) classifier on the samples from the source distribution with re-labeled source classes (matched with the target classes), (4) uses the samples of the matched source classes to re-train the target classifier along with Directional Regularization that builds the final model for the target data. This entire procedure is pictorially depicted in Figure 2.

## Implementation

The idea of GWS detailed in the previous section assumes that the source and the target distributions are ‘similar’ under a certain feature transformation. The raw video data in the RGB space does not adhere to such assumptions because of the variability of the content/scene/subjects etc. However, transformations such as optical flow masks most of the non-motion related information and exaggerates the motion information. In this domain, it is reasonable to assume that the source and target action classes are ‘similar’ (Refer Figure 1). We match each target Autism class to a source class using the baseline Autism model  $\mathbf{P}_\theta^t(\mathbf{y}_t|\mathbf{x}_t)$ . Since a probabilistic Softmax layer is used at the output of the classifier,

the matched source class, for every given target class, can be simply taken to be that source class whose samples get labeled as the given target class most of the times as compared to all other source classes. The optimization problem in Eq. 4 can be approximated as follows:

$$\mathbf{y}_s^* | \mathbf{y}_t^q = \operatorname{argmax}_{\mathcal{Y}_s} \mathcal{L}_{\mathbf{y}_t = \mathbf{y}_t^q | \mathbf{x}_s} \quad (8)$$

$$= \prod_{j=1}^1 \mathbf{P}_\theta^t(\mathbf{y}_{tj} = \mathbf{y}_t^q | \mathbf{x}_{sj}) \quad (9)$$

$$\approx \operatorname{argmax}_{\mathcal{Y}_s} \left[ \operatorname{Count}_{j \in \{1, \dots, 1\}} \left( \mathbf{y}_t^q = \operatorname{argmax}_{\mathcal{Y}_t} \mathbf{P}_\theta^t(\mathbf{y}_{tj} | \mathbf{x}_{sj}) \right) \right] \quad (10)$$

We propose to use state-of-the-art action recognition models such as I3D and TSN that employ optical flow streams. Baseline target Autism model  $\mathbf{P}_\theta^t(\mathbf{y}_t|\mathbf{x}_t)$  is obtained by training Autism data using an I3D or TSN architecture by initializing their weights with pre-trained Kinetics/ImageNet/HMDB51 models. GWS and re-training with Directional Regularization are performed only on the flow stream without altering the RGB stream. Once the source classes are matched, the samples from the matched source classes are trimmed to a fixed length of average duration of the target clips ( $\approx 2$  seconds in the Autism data). This is accomplished by trimming each video sample of the matched source classes using an overlapping window approach and taking that portion of the video which has the highest Softmax prediction score under the baseline target classifier for that particular target class. These trimmed (or action localized) video clips from the source data are used for data augmentation to re-train the target classifier with the Directional Regularization loss. Both the RGB and re-trained optical flow streams are combined for final prediction. All the data is pre-processed to generate the optical flow frames using the TVL1 algorithm (Zach, Pock, and Bischof 2007). The code associated with this paper can be found at <https://github.com/prinshul/GWSDR>.

## Autism Dataset

The target Autism data consists of 37 subjects (‘child’ and ‘subject’ are used interchangeably). Five subjects have Caucasian origin and the rest are Asians. Their ages range from 2-14. During assessment, a clinician performs the functional assessment by probing a child for age-appropriate listener response and imitation skills by invoking an instruction response and expecting a child to respond through a human action. We deliberately chose eight representative human action responses invoked through either listener response or imitation instruction for our experiments. Specifically, the action classes selected are - ‘Move the table’ and ‘Arms up’ for gross motor skill assessment, ‘Lock hands’ and ‘Tapping’ for fine motor skills, ‘Rolly polly’ for complex motor action, ‘Touch nose’, ‘Touch head’ and ‘Touch ear’ for identification of different parts of the body. We chose these particular actions since the presence of age-appropriate fine and gross motor skills demonstrate neuro-typical develop-

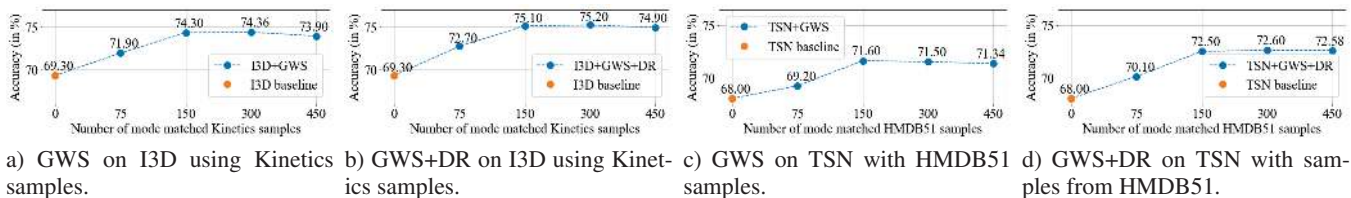


Figure 3: Performance on I3D and TSN with different baselines and different amount of augmented source (Kinetics and HMDB51) samples using GWS and DR.

ment of a child, as well as providing a clear picture of atypical development (Gowen and Hamilton 2013). A total of 1481 video clips were recorded in a semi-structured environment with the clinician facing the child, and three synchronized cameras (Logitech C922 720p 60fps) were placed to record the videos. The first camera faced the clinician, the second faced the child and the third was positioned laterally to both the clinician and the child. Figure 4 depicts representative frames of some action classes during training sessions for different camera positions. The video clips were annotated by trained clinical psychologists. The response of the child for a particular stimulus is treated as a human action response classification problem. 2D-videos were recorded with 37 subjects and 12 licensed clinicians in two phases for a period of 12 months in the US and Asia across different Autism diagnostic centers. There are 148 annotated video clips and the test set contains 1333 clips. All the training and test samples are trimmed video clips containing exactly one of the eight actions of 2-5 seconds in duration. A copy of the Autism dataset can be obtained by contacting the authors.

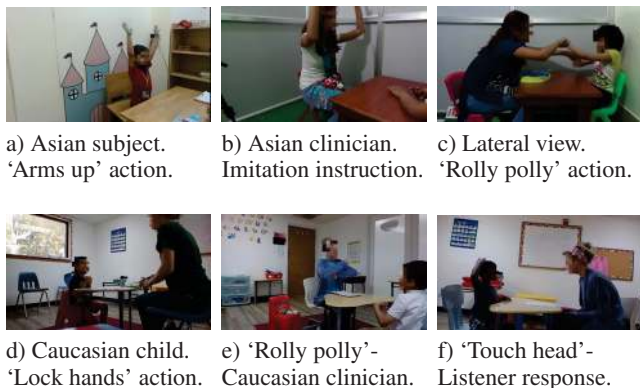


Figure 4: Training sessions of Asian and Caucasian subjects by clinicians with different camera positions.

## Experiments and Results

Two large-scale publicly available human action recognition datasets namely Kinetics (Kay et al. 2017) and HMDB51 (Kuehne et al. 2011) are used as the source datasets while the data described in the previous section (termed as the Autism dataset) is the target dataset. The task is of standard 8-class classification on the target Autism data with classes as described in the previous section. We train a bi-

nary action classifier with two classes, ‘Response’ and ‘No Response’. The ‘Response’ class has samples from all the 8 action classes while the ‘No Response’ class has samples containing random actions (different from the 8 actions in the Autism dataset) or no action at all. This classifier helps to mark the video clips with actions not belonging to one of the 8 classes or having no response. All the further experiments are carried out on video clips containing a ‘Response’.

### Baselines

For mode matching experiments, I3D is used in conjunction with Kinetics and TSN with HMDB51. Our baselines are I3D and TSN which are pre-trained on Kinetics and HMDB51 respectively. Next, using the baseline model, the Kinetics and HMDB51 classes are mode matched to Autism classes. Table 1 shows matched classes on both the source datasets. It is apparent from Table 1 that mode matching maps semantically similar actions from Kinetics and HMDB51 to the Autism actions, confirming the proposed hypothesis. Figure 3 shows the variation of test ac-

Autism	Kinetics	HMDB51
Move the table	Pushing cart	Push
Touch ear	Tying necktie	Sit up
Lock hands	Playing trombone	Shake hands
Touch head	Blowdrying hair	Shoot ball
Touch nose	Putting on eyeliner	Eat
Rolly polly	Playing hand clapping games	Flic flac
Tapping	Playing drums	Chew
Arms up	Jumping jacks	Fall floor

Table 1: Action classes from the source datasets (Kinetics and HMDB51) matched to target (Autism) classes.

curacy with different amounts of source data from Kinetics and HMDB51 mode matched classes augmented with Autism classes. Baseline models are re-trained with varying amounts of mode matched Kinetics or HMDB51 samples (equivalent to 5% ( $\approx 75$  samples) to 30% ( $\approx 450$  samples) of all the Autism samples). In all of these baseline models, when the size of augmented source data is increased and the model is re-trained, the baseline accuracy increases till the percentage of the augmented source data is comparable (in terms of number of samples) with the Autism data. It is seen that when the source data dominates the Autism data, the accuracy drops. This is expected because when the source distribution dominates, the classifier tends to overfit

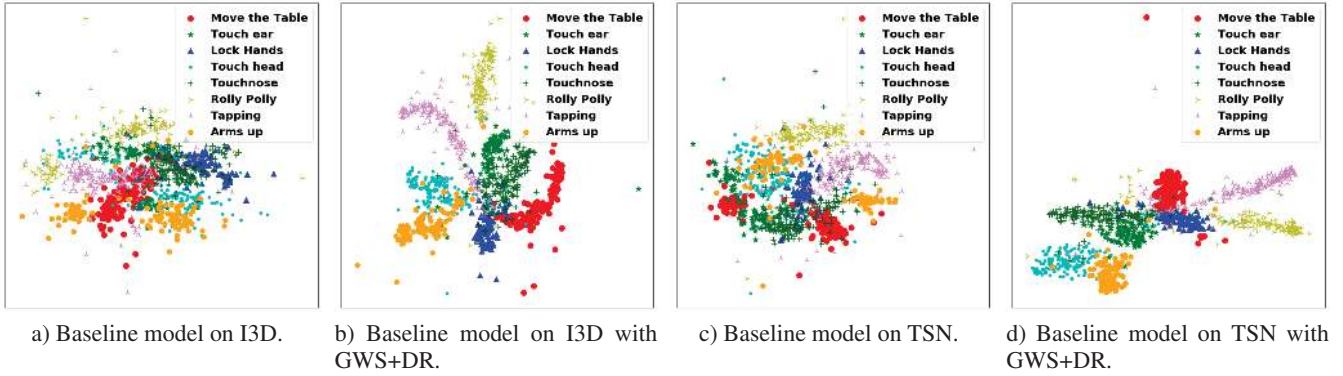


Figure 5: t-SNE plots of embeddings of penultimate layers of baseline I3D and TSN Autism models with and without GWS and DR. It is clearly seen that the inter-class separability has increased and clusters are more dense after GWS and DR.

on it. However, with GWS along with DR, not only the rise in test accuracy is more, the drop in accuracy after peaking is smoother as compared to GWS without DR. This implies that DR offers more tolerance towards the augmented source data which allows the performance to increase further. Our approach outperforms all baseline Autism models of I3D and TSN with comparable source samples. When we re-train the baseline Autism I3D classifier by augmenting random samples from the source (without mode matching), the test accuracy drops to 32% thereby showing importance of GWS. t-SNE (Maaten and Hinton 2008) plots for the 8-dimensional embeddings from penultimate layer of I3D and TSN are obtained for the baseline models with and without GWS and DR as shown in Figure 5. With our approach, the inter-class separability of samples has increased while the intra-class separability has decreased so the model tends to be more confident in its predictability of Autism classes.

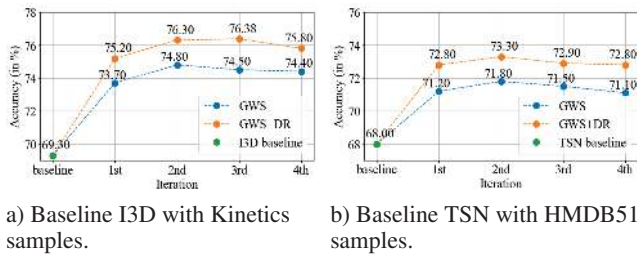


Figure 6: Performance on I3D and TSN with iteration over the source samples with GWS and DR. The accuracy increases through iteration by augmenting with newer mode matches samples in every iteration although overfitting on source data occurs from 3<sup>rd</sup> iteration.

In the next set of experiments, we iteratively re-train the target model with GWS and DR. That is, in every new iteration we discard the mode matched source samples from the previous iteration keeping the number of samples similar in every iteration (equivalent to the number of Autism training samples) and re-train the target model with new set of samples from the matched modes classes. Figure 6 shows the variation in accuracy when the new source samples are

Model	Asian	Caucasian	≤ 5 yrs.	> 5 yrs.
I3D	58.2%	34.6%	55.1%	41.2%
TSN	53.7%	30.2%	51.7%	38.1%
I3D+GWS+DR	<b>63.7%</b>	<b>40.2%</b>	<b>59.8%</b>	<b>46.5%</b>
TSN+GWS+DR	<b>57.0%</b>	<b>33.8%</b>	<b>55.4%</b>	<b>40.9%</b>

Table 2: Performance of GWS and DR with a specific bias in the Autism training dataset. It is seen that in all the cases, the proposed approach offers better performance over the baselines.

augmented with Autism data in every iteration. In Figure 6a, the test accuracy increases from the baseline for initial iterations with GWS and DR. With GWS only (I3D+GWS), the accuracy drops after the 2<sup>nd</sup> iteration. If GWS is applied with DR on the baseline (I3D+GWS+DR), it continues to increase even after the 2<sup>nd</sup> iteration but starts to dip from 3<sup>rd</sup> iteration on-wards which can be ascribed to overfitting on the source data. As with the previous cases, we observe more tolerance of DR with newer samples as compared to GWS. Newer samples are accepted with lesser surprise in DR which enhances the generalizability and performance. Similar behavior is observed with TSN as shown in Figure 6b.

### Bias in the Training Data

Table 2 shows the results for the proposed method under different kinds of dataset biases. The results in the second column in this table are accuracies when the training data has only Asian subjects. Third column are test accuracies when the training data has Caucasian subjects. In the fourth column, the training data has subjects that are 5 years or below. The last column are test accuracies when the training data has subjects above 5 years in age. It is observed that even if the dataset has a bias, our approach performs better than the baselines.

### GWS and DR under Different Settings

Table 3 shows test accuracy scores on baseline Autism model under six different settings. We sampled examples



Model	Second best modes	Flow + RGB stream	Handpicked modes	Combined modes	I3D→TSN	TSN→I3D
I3D+GWS	<b>73.5%</b>	67.5%	<b>74.2%</b>	<b>73.8%</b>	-	<b>73.2%</b>
TSN+GWS	<b>71.4%</b>	65.8%	<b>71.9%</b>	<b>71.2%</b>	<b>71.7%</b>	-
I3D+GWS+DR	<b>74.4%</b>	68.1%	<b>75.3%</b>	<b>74.8%</b>	-	<b>74.4%</b>
TSN+GWS+DR	<b>72.4%</b>	66.7%	<b>72.8%</b>	<b>72.7%</b>	<b>73.1%</b>	-

Table 3: GWS and DR under different settings on baselines (69% for I3D and 68% for TSN) - It is seen that (a) GWS or GWS+DR with second best modes too leads in better performance (b) re-training the RGB stream is detrimental since there is no similarity in the RGB space (c) GWS or GWS+DR with hand-picked modes also results in improvement in accuracy (d) re-training with samples from matched modes classes from different datasets results in performance enhancement (e) & (f) Cross neural architecture GWS and DR - Increase in accuracy from the baselines indicate that similar actions in the optical flow space retain their meaning irrespective of the neural architectures.

Model	Autism Dataset	SSBD
TSN (Pretrained with HMDB51) (Wang et al. 2016)	68.0%	87.4%
I3D (Pretrained with Kinetics) (Carreira and Zisserman 2017)	69.3%	91.2%
ECO (Pretrained with Kinetics) (Zolfaghari, Singh, and Brox 2018)	61.4%	80.1%
TSM (Pretrained with Kinetics) (Lin, Gan, and Han 2019)	69.8%	90.5%
R(2+1)D (Pretrained with IG-65M) (Ghadiyaram, Tran, and Mahajan 2019)	68.4%	88.3%
TSN+DR	70.1%	89.2%
I3D+DR	71.3%	92.8%
TSN+GWS	71.6%	90.3%
I3D+GWS	74.3%	93.6%
TSN+GWS+DR	72.5%	91.4%
I3D+GWS+DR	<b>75.1%</b>	<b>95.7%</b>

Table 4: Comparison with state-of-the-art action recognition models on Autism dataset and SSBD.

from second best modes to augment with Autism data using GWS and DR. The results in the second column are the test accuracy scores using second best modes. It can be seen that the performance is better than the baseline model albeit less than the performance of model with the best modes shown in Table 1. It is however apparent that even second best modes preserve the closeness in the optical flow space. All the experiments are executed by re-training optical flow stream of the baseline models. Besides optical flow, if we re-train the RGB stream as well, it is seen that the performance of the classifiers deteriorates as shown in third column of Table 3. This ascertains the fact that the notion of closeness is valid only for optical flow samples. Next, we handpicked similar action classes (like ‘Washing hands’ in Kinetics matched to ‘Lock hands’ in Autism data, etc.) from Kinetics and HMDB51 datasets and applied our approach on the baseline models. Fourth column of Table 3 records test accuracy scores with these handpicked classes (or modes). The results are comparable to our approach where we employ baseline models to find similar classes or actions using mode matching. Hence, irrespective of the metric used to find similarity (either using human intelligence or mode matching), the performance is much better than the baselines when they are augmented and re-trained with similar classes (GWS) with or without DR. Next, we combined corresponding best mode matched classes from Kinetics and HMDB51 (like ‘Pushing cart’ from Kinetics is combined with ‘Push’

from HMDB51, etc.) and used these samples for GWS and DR as shown in the fifth column of Table 3. The performance is comparable with our approach when the modes are not combined which implies that the idea of similarity in optical flow space is preserved even if different datasets are merged. Fifth and sixth columns report test accuracies for the case when the modes matched using one architecture is used to augment the classifier built on another architecture. Fifth column of Table 3 are the test accuracy scores when baseline TSN is re-trained with augmented Kinetics modes extracted from mode matching on I3D. Similarly, sixth column has test accuracy scores when baseline I3D is re-trained with HMDB51 modes extracted from TSN. The scores are still better than the baseline Autism model (69% for I3D and 68% for TSN) which indicates that similar actions in the optical flow space retain their meaning irrespective of the neural architectures used as a backbone.

### Comparisons with State-of-the-art

Table 4 reports test accuracy for GWS and DR on baseline I3D and TSN. It is consistently observed that with our approach (GWS or DR or GWS+DR), the performance is better than the baseline I3D and TSN and state-of-the-art action recognition models using Autism dataset and SSBD. SSBD is a public dataset having 75 examples and three action classes namely ‘arm flapping’, ‘head banging’ and ‘spinning’ that are used in Autism diagnosis. Additionally, it is

seen that DR offers additional accuracy benefits over GWS.

## Conclusion

In this paper, we proposed a method for improving the generalization abilities of a classifier designed for human action recognition trained on scarce data. Specifically, leveraging the semantic similarities of the action classes in the optical flow space, we proposed a generic method called Guided Weak Supervision (GWS) to augment and re-train a classifier on the target data with samples from a large-scale annotated dataset, along with a novel loss function termed as Directional Regularization (DR) that would result in performance enhancement. We demonstrated the efficacy of the proposed method for screening, diagnosis and behavioral treatment for ASD by treating imitation, listener response and motor skills as action classification tasks. With the proposed framework, we can integrate and automate complete value chain of screening to treatment for children with ASD by capturing behavioral treatment progress data temporally and remotely. The other future direction is to employ the techniques of GWS and DR for generalized supervised learning tasks beyond the proposed use case of Autism.

## References

- Carreira, J., and 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*, 6299–6308.
2019. Data & statistics on autism spectrum disorder. <https://www.cdc.gov/ncbddd/autism/data.html>. Accessed: 2019-06-30.
- Donahue, J.; Anne Hendricks, L.; Guadarrama, S.; Rohrbach, M.; Venugopalan, S.; Saenko, K.; and Darrell, T. 2015. Long-term recurrent convolutional networks for visual recognition and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2625–2634.
- Estes, A.; Munson, J.; Rogers, S. J.; Greenson, J.; Winter, J.; and Dawson, G. 2015. Long-term outcomes of early intervention in 6-year-old children with autism spectrum disorder. *Journal of the American Academy of Child & Adolescent Psychiatry* 54(7):580–587.
- Ghadiyaram, D.; Tran, D.; and Mahajan, D. 2019. Large-scale weakly-supervised pre-training for video action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 12046–12055.
- Gowen, E., and Hamilton, A. 2013. Motor abilities in autism: a review using a computational context. *Journal of autism and developmental disorders* 43(2):323–344.
- Horn, B. K., and Schunck, B. G. 1981. Determining optical flow. *Artificial intelligence* 17(1-3):185–203.
- Kay, W.; Carreira, J.; Simonyan, K.; Zhang, B.; Hillier, C.; Vijayanarasimhan, S.; Viola, F.; Green, T.; Back, T.; Natsev, P.; et al. 2017. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*.
- Kuehne, H.; Jhuang, H.; Garrote, E.; Poggio, T.; and Serre, T. 2011. Hmdb: a large video database for human motion recognition. In *2011 International Conference on Computer Vision*, 2556–2563. IEEE.
- Lin, J.; Gan, C.; and Han, S. 2019. Tsm: Temporal shift module for efficient video understanding. In *Proceedings of the IEEE International Conference on Computer Vision*, 7083–7093.
- Maaten, L. v. d., and Hinton, G. 2008. Visualizing data using t-sne. *Journal of machine learning research* 9(Nov):2579–2605.
- Mandal, D.; Narayan, S.; Dwivedi, S. K.; Gupta, V.; Ahmed, S.; Khan, F. S.; and Shao, L. 2019. Out-of-distribution detection for generalized zero-shot action recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Rapin, I. 1991. Autistic children: Diagnosis and clinical features. *Pediatrics* 87(5):751–760.
- Simonyan, K., and Zisserman, A. 2014. Two-stream convolutional networks for action recognition in videos. In *Advances in neural information processing systems*, 568–576.
- Soomro, K.; Zamir, A. R.; and Shah, M. 2012. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*.
- Sundar Rajagopalan, S.; Dhall, A.; and Goecke, R. 2013. Self-stimulatory behaviours in the wild for autism diagnosis. In *The IEEE International Conference on Computer Vision (ICCV) Workshops*.
- Tran, D.; Bourdev, L.; Fergus, R.; Torresani, L.; and Paluri, M. 2015. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, 4489–4497.
- Wang, L.; Xiong, Y.; Wang, Z.; Qiao, Y.; Lin, D.; Tang, X.; and Van Gool, L. 2016. Temporal segment networks: Towards good practices for deep action recognition. In *European conference on computer vision*, 20–36. Springer.
- Wang, Y.; Zhou, L.; and Qiao, Y. 2018. Temporal hallucinating for action recognition with few still images. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Yang, H.; He, X.; and Porikli, F. 2018. One-shot action localization by learning sequence matching network. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Zach, C.; Pock, T.; and Bischof, H. 2007. A duality based approach for realtime tv-l 1 optical flow. In *Joint pattern recognition symposium*, 214–223. Springer.
- Zhu, L., and Yang, Y. 2018. Compound memory networks for few-shot video classification. In *The European Conference on Computer Vision (ECCV)*.
- Zolfaghari, M.; Singh, K.; and Brox, T. 2018. Eco: Efficient convolutional network for online video understanding. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 695–712.