

Efficient Multiple Instance Convolutional Neural Networks for Gigapixel Resolution Image Classification

Le Hou¹, Dimitris Samaras¹, Tahsin M. Kurc², Yi Gao^{2,1,3}, James E. Davis⁴, and Joel H. Saltz^{2,1,4,5}

¹Department of Computer Science, Stony Brook University, NY, USA
{lehou,samaras}@cs.stonybrook.edu

²Department of Biomedical Informatics, Stony Brook University, NY, USA
{tahsin.kurc,joel.saltz}@stonybrook.edu

³Department of Applied Mathematics and Statistics, NY, USA

⁴Department of Pathology, Stony Brook Hospital, NY, USA
{yi.gao,james.davis}@stonybrookmedicine.edu

⁵Cancer Center, Stony Brook Hospital, NY, USA

Abstract

Convolutional Neural Networks (CNNs) are state-of-the-art models for many image and video classification tasks. However, training on large-size training samples is currently computationally impossible. Hence when the training data is multi-gigapixel images, only small patches of the original images can be used as training input. Since there is no guarantee that each patch is discriminative, we advocate the use of Multiple Instance Learning (MIL) to combine evidence from multiple patches sampled from the same image. In this paper we propose a framework that integrates MIL with CNNs. In our algorithm, patches of the images or videos are treated as instances, where only the image- or video-level label is given. Our algorithm iteratively identifies discriminative patches in a high resolution image and trains a CNN on them. In the test phase, instead of using voting to the predict the label of the image, we train a logistic regression model to aggregate the patch-level predictions. Our method selects discriminative patches more robustly through the use of Gaussian smoothing. We apply our method to glioma (the most common brain cancer) subtype classification based on multi-gigapixel whole slide images (WSI) from The Cancer Genome Atlas (TCGA) dataset. We can classify Glioblastoma (GBM) and Low-Grade Glioma (LGG) with an accuracy of 97%. Furthermore, for the first time, we attempt to classify the three most common subtypes of LGG, a much more challenging task. We achieved an accuracy of 57.1% which is similar to the inter-observer agreement between experienced pathologists.

1. Introduction

For many image and video classification tasks, Convolutional Neural Networks (CNNs) are currently the state-of-the-art classifiers [22, 21, 5]. However, due to high computational cost, CNNs cannot be applied to very high resolution images, such as Whole Slide Tissue Images (WSI). We are motivated by the problem of WSI-based automatic glioma classification. Gliomas are one of the most devastating types of tumors. They are complex and highly fatal. Better diagnosis and classification into grades and subtypes of these tumors is critical to the study of disease onset and progression as well as the development of targeted therapies. The effects of cancer show as changes in tissue at the cellular and sub-cellular scales. Thus, high resolution (gigapixel) microscopy images of tissue slides provide rich information with which to study gliomas. We need to train a CNN model on patches extracted from the full-size image. However, the ground truth labels of individual patches are unknown, as only the image-level ground truth label is given. Classification of gliomas into grades and subtypes using tissue slides and microscopy images is a challenging problem, because tumors are heterogeneous and may have a mixture of structures and texture properties. In addition, morphological indicators of transitions from one subtype to another can be subtle. For these reasons, patch-level classifications are not necessarily consistent with the WSI-level classification.

We propose using Multiple Instance Learning (MIL) as part of a two-level model, shown in Fig. 1. The first-level model is an Expectation Maximization (EM) based MIL method combined into a CNN that outputs patch-level predictions. In particular, we assume that there is a hidden

variable associated with each patch extracted from an image that indicates whether the patch is discriminative or not. Here, "a patch being discriminative" means that the true hidden label of the patch is the same as the true label of the image. Initially, we consider all patches to be discriminative and train the CNN model that outputs the cancer type probability for each input patch. We apply a spatial smoothing to the resulting probability map and select only patches with higher probability values as discriminative patches. We iterate this process using the new set of discriminative patches in an EM fashion until convergence. In the second-level, the histogram of patch-level predictions is input into a multi-class logistic regression that predicts the image-level label.

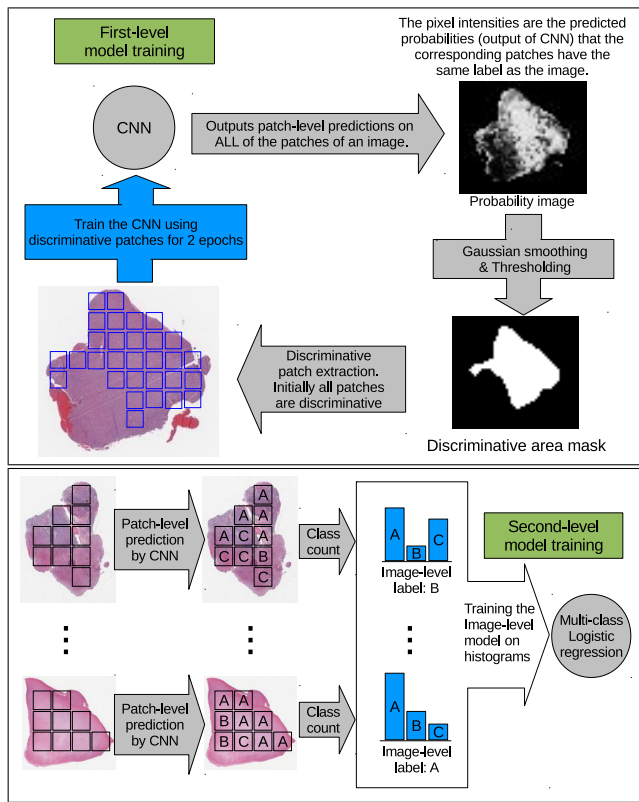


Figure 1: An overview of our workflow of glioma pathology image classification (best viewed in color). The assumption is that discriminative information is encoded in high resolution details and some of the patches extracted from the image do not contain such discriminative information. Top: An EM-based MIL-CNN that iteratively identifies non-discriminative patches and eliminates them from the training set. Bottom: The second-level model learns how to aggregate patch-level predictions to the image-level prediction.

Pathology image classification and segmentation is an

active field of research and development. There are three major categories of algorithms and corresponding applications: Patch-level classification and/or segmentation methods [43, 8, 2, 34, 11] require that a lot of patch labels are provided. Certain methods [14, 25] integrate patch-level information into supervised WSI classification. One method [25] uses a domain-specific pipeline for WSI-based Glioblastoma (GBM) vs. Low-Grade Glioma (LGG) classification and detects discriminative regions such as necrotic and microvascular proliferation areas are detected via a fully supervised model. This method then applies a decision tree for the final prediction of classes. The domain-specific discriminative features do not generalize to other cancer types, not even to subtype glioma classification. The need to integrate patch-level information without cancer type specific features led to MIL-based WSI classification [13, 37, 38, 19].

In the MIL paradigm [15, 24, 3], unlabeled instances belong to labeled bags of instances. The goal is to predict the label of a new bag and/or the label of each instance. The Standard Multi-Instance (SMI) assumption [15] states that for a binary classification problem, the bag is positive if and only if there exists at least one positive instance in the bag. This led to the MIL version of different classifiers [4, 41, 20], including MIL with Neural Networks (NN) [32, 44, 23]. The Back Propagation for Multi-Instance Problems (BP-MIP) [32, 44] performs back propagation along the instance with the maximum response if the bag is positive. This is inefficient because only one instance per bag is trained in one training iteration on the whole bag. An efficient neural network with MIL is proposed [23] with the drawback that it only works on the feature space due to the need of averaging over the instances.

In image classification, the image is the bag and image-windows are the instances [26]. Thus, MIL-based CNNs are applied to object recognition [28] and semantic segmentation [30]. In these methods, the training error is only propagated through the object-containing window which is also assumed to be the window that has maximum prediction confidence. We argue that this is not robust since training only on high-confidence windows may lead to overfitting. In recent semantic image segmentation approaches [9, 31, 29], a smoothing algorithm is applied on the output probability (feature) maps of CNNs. In this way, irrelevant windows can be identified more robustly.

We generalized the previous idea and extended it to the patch-based image classification problem based on two assumptions: First, in each image, there may exist some non-discriminative patches. Their true latent labels do not match the label of the image. Second, non-discriminative patches also contribute to the image-level labels jointly. For example [3], a beach image may contain a sand-patch and a water-patch; neither of them are by themselves beach

patches. Based on these two assumptions, our method consists of two levels. The first-level model efficiently eliminates non-discriminative patches from the final CNN training whereas the second-level logistic regression predicts the image-level label based on the patch-level predictions [10]. The first layer model reduces the amount of noise and therefore improves performance. The second-level model models a more general multiple instance assumption [16] which fits the assumption of our main application of WSI-based glioma classification. It also can be interpreted as a stacked generalization method [36].

Our main contributions are as follows.

1. We propose a new classification model that combines MIL with CNNs that work in raw pixel space. It utilizes the spatial relationship between patches extracted from image to identify discriminative patches. Additionally, our algorithm adds very little computational cost to the conventional supervised CNNs.
2. We propose a second-level model that can be applied on datasets that do not meet the SMI assumption. On the glioma classification dataset, this method outperforms the first-level model with the SMI assumption significantly.
3. Our model achieves state-of-the-art results classifying Glioblastoma (GBM) and Low-Grade Glioma (LGG) on the TCGA dataset. Furthermore, for the first time, we can classify three major subtypes of LGG with an accuracy similar to inter-observer agreement between experienced pathologists.

The rest of this paper is organized as follows. Sec. 2 describes the framework of the EM-based MIL algorithm. Sec. 3 discusses the identification of discriminative patches. Sec. 4 explains the second-level model that predicts the image-level label by aggregating patch-level predictions. Sec. 5 shows experimental results. The paper concludes in Sec. 6. App. A gives glioma subtype details.

2. EM-based MIL with CNN

An overview of our EM-based method can be found in Fig. 1. We model the high resolution image or video as a bag and the patches extracted from it as instances. We have a ground truth label for the whole image but not for the individual patches. Assuming some instances are not discriminative, *i.e.* the true label of the instances does not match the label of the bag, whether an instance is discriminative or not is modeled as a hidden binary variable.

We denote $X = \{X_1, X_2, \dots, X_N\}$ as the dataset containing N bags. Each bag $X_i = \{X_{i,1}, X_{i,2}, \dots, X_{i,N_i}\}$ consists of N_i instances, where $X_{i,j} = \langle x_{i,j}, y_i \rangle$ is the j -th instance and its associated label in the i -th bag. Assuming

the bags are independent and identically distributed (i.i.d.), the X and the hidden variables H are generated by the following generative model:

$$\begin{aligned} P(X, H) &= \prod_{i=1}^N P(X_i | H_i) P(H_i) \\ &= \prod_{i=1}^N \left(P(X_{i,1}, X_{i,2}, \dots, X_{i,N_i} | H_i) P(H_i) \right), \end{aligned} \quad (1)$$

where the hidden variable $H = \{H_1, H_2, \dots, H_N\}$, where $H_i = \{H_{i,1}, H_{i,2}, \dots, H_{i,N_i}\}$ and $H_{i,j}$ is the hidden variable that indicates whether instance $x_{i,j}$ is discriminative for label y_i of bag X_i . We further assume that all $X_{i,j}$ depends on $H_{i,j}$ only and are independent with each other given $H_{i,j}$. Thus

$$P(X, H) = \prod_{i=1}^N \prod_{j=1}^{N_i} \left(P(X_{i,j} | H_{i,j}) P(H_{i,j}) \right). \quad (2)$$

We maximize the data likelihood $P(X)$ using the EM algorithm.

1. At the initial E step, we set $H_{i,j} = 1$ for all i, j . This means that all instances are considered discriminative.
2. M step: We update the model parameter θ to maximize the data likelihood

$$\begin{aligned} \theta &\leftarrow \arg \max_{\theta} P(X | H; \theta) \\ &= \arg \max_{\theta} \prod_{x_{i,j} \in D} P(x_{i,j}, y_i | \theta) \\ &\quad \times \prod_{x_{p,q} \notin D} P(x_{p,q}, y_q | \theta), \end{aligned} \quad (3)$$

where D is the set of discriminative patches. We assume a uniform generative model for all non-discriminative instances. Therefore the optimization problem in Eq. 3 can be simplified to

$$\begin{aligned} &\arg \max_{\theta} \prod_{x_{i,j} \in D} P(x_{i,j}, y_i | \theta) \\ &= \arg \max_{\theta} \prod_{x_{i,j} \in D} P(y_i | x_{i,j}; \theta) P(x_{i,j}). \end{aligned} \quad (4)$$

Additionally we assume a uniform distribution over $x_{i,j}$. Thus Eq. 4 describes a discriminative model (in this paper we use a CNN).

- E step: We estimate the hidden variables H . In particular, $H_{i,j} = 1$ if and only if $P(H_{i,j} | X)$ is above a certain threshold. In the case of image classification, given the i -th image, $P(H_{i,j} | X)$ is obtained by applying Gaussian smoothing on $P(y_i | x_{i,j}; \theta)$ (Detailed in Sec 3). This smoothing step utilizes the spatial relationship of $P(y_i | x_{i,j}; \theta)$ in the image. We then iterate back to the M step till convergence.

Many MIL algorithms can be interpreted as a version of this formulation. Based on the SMI assumption, the instance with the maximum $P(H_{i,j} | X)$ is considered as the discriminative instance for the positive bag, as in the EM Diverse Density (EM-DD) [42] and the BP-MIP [32, 44] algorithms. In the latter case, a single gradient descent is performed in the M step per positive bag, which makes the algorithm computationally inefficient. In the field of semantic image segmentation, [29] uses a fully connected Conditional Random Field (CRF) to model $P(H | X)$ and then perform the relevant EM-based updates.

3. Discriminative instance selection

This section focuses on the E step of estimating $P(H | X)$ and the way to choose a threshold for determining whether a patch is discriminative. After this step, the patches $x_{i,j}$ that have $P(H_{i,j} | X)$ larger than a threshold $T_{i,j}$ are considered as discriminative and are selected to continue training the CNN.

To estimate $P(H_{i,j} | X)$, it is straightforward to assume that $P(H_{i,j} | X)$ is correlated with $P(y_i | x_{i,j}; \theta)$, *i.e.* patches with lower $P(y_i | x_{i,j}; \theta)$ tend to have lower probability $x_{i,j}$ to be discriminative. However, a hard-to-classify patch, or a patch close to the decision boundary may have low $P(y_i | x_{i,j}; \theta)$ as well. These patches are informative and should not be rejected as non-discriminative. Therefore, to obtain a more robust $P(H_{i,j} | X)$, we apply the following two steps: First, we train two CNNs on two different scales in parallel and the prediction $P(y_i | x_{i,j}; \theta)$ is the averaged prediction of the two CNNs. Second, we simply denoise the probability map $P(y_i | x_{i,j}; \theta)$ of each image with a Gaussian kernel to obtain $P(H_{i,j} | X)$.

Choosing a thresholding scheme carefully yields significantly better performance than a simpler thresholding scheme [29]. We obtain the threshold $T_{i,j}$ for $P(H_{i,j} | X)$ as follows: We note S_i as the set of $P(H_{i,j} | X)$ values for all $x_{i,j}$ of the i -th image and E_c as the set of $P(H_{i,j} | X)$ values for all $x_{i,j}$ of the c -th class. We introduce the image-level threshold H_i as the P_1 -th percentile of S_i and the class-level threshold R_i as the P_2 -th percentile of E_c , where P_1 and P_2 are predefined parameters. Then the threshold $T_{i,j}$ is defined as the minimum value between H_i and R_i . There are two advantages of our method. First, by using the image-level threshold, there are at least $1 - P_1$ percent of

patches that are considered discriminative for each image. Second, by using the class-level threshold, the thresholds can be easily adapted to classes with different prior probabilities.

Layer	Filter size, stride	Output size
Input	-	$400 \times 400 \times 3$
Conv	$10 \times 10, 2$	$196 \times 196 \times 80$
ReLU+LRN	-	$196 \times 196 \times 80$
Max-pool	$6 \times 6, 4$	$49 \times 49 \times 80$
Conv	$5 \times 5, 1$	$45 \times 45 \times 120$
ReLU+LRN	-	$45 \times 45 \times 120$
Max-pool	$3 \times 3, 2$	$22 \times 22 \times 120$
Conv	$3 \times 3, 1$	$20 \times 20 \times 160$
ReLU	-	$20 \times 20 \times 160$
Conv	$3 \times 3, 1$	$18 \times 18 \times 200$
ReLU	-	$18 \times 18 \times 200$
Max-pool	$3 \times 3, 2$	$9 \times 9 \times 200$
FC	-	320
ReLU+Drop	-	320
FC	-	320
ReLU+Drop	-	320
FC	-	6
Softmax	-	6

Table 1: The architecture of our CNN used in glioma classification. The output size is given in width \times height \times the number of feature maps. ReLU+LRN is a sequence of Rectified Linear Units (ReLU) followed by Local Response Normalization (LRN). Similarly, ReLU+Drop is a sequence of ReLU followed by dropout. The dropout probability is 0.5.

4. Image label prediction

In this section we describe how we combine the patch-level classifiers of Sec. 3 to predict accurately the image-level label. We input all the predictions of the patch-level classifiers into a multi-class logistic regression that predicts the image-level label. There are two reasons for combining multiple instances: First, on difficult datasets, we do not want to assign an image-level prediction just based simply based on a single patch-level prediction (as is the case of the SMI assumption [15]). Second, even though certain patches are not discriminative individually, their joint appearance might be discriminative. For example, Oligoastrocytoma (OA) (see App. A) is a "mixed" type glioma that is diagnosed when two single glioma types (Oligodendroglioma and Astrocytoma, possibly on different patches) are jointly present on the slide.

In particular, the class histogram of the patch-level predictions is the input to our linear multi-class logistic regression model [6]. This model uses an L2 regularizer which is

selected by cross-validation on a validation dataset. A linear model is sufficient in this case because the inputs of the model are expected to be directly correlated with the output labels. This model can be thought of a count-based multiple instance learning method with two-level learning [35].

5. Experiments

We have evaluated our method on two datasets: The Cancer Genome Atlas (TCGA) dataset [1] for Whole Slide Image (WSI) based glioma classification and the Kinect interaction dataset [39] for two-person interaction classification.

5.1. WSI-based glioma classification

We have evaluated our algorithm to the glioma classification problem on the TCGA dataset. Based on large scale experiments we show that our model outperforms existing methods significantly.

5.1.1 TCGA dataset

The TCGA dataset contains detailed clinical information and the Hematoxylin and Eosin (H&E) stained WSIs of 6 types of gliomas. A typical resolution of a WSI is 100k by 50k pixels. Fig. 2 shows sample patches of each class. The numbers of WSIs and patients in each class are shown in Tab. 2. All classes are described in the App. A.

Glioma	# of WSIs	# of patients
GBM	510	209
OD	206	100
OA	183	106
DA	114	82
AA	36	29
AO	15	13

Table 2: The numbers of WSIs and patients in each class from the TCGA dataset. For descriptions of each class, see the App. A.

5.1.2 Patch extraction and augmentation

To train the CNN model, patches of size 500 by 500 are extracted from the WSI. To capture structures in different scales, we extract patches from two scales: 20X (0.5 microns per pixel) and 5X (2.0 microns per pixel). Patches that contain less than 30% tissue sections or have too much blood are discarded. For each WSI, patch stride is computed carefully to make sure around 1000 valid patches per image per scale are extracted. Therefore, in most cases the patches are non-overlapping given the resolution of the WSI.

To prevent the CNN from severe overfitting, we perform three kinds of data augmentation. First, a random S by S

sub-patch is selected, where S is a predefined hyperparameter. Second, the sub-patch is randomly rotated and mirrored. Third, the amount of Hematoxylin and eosin stained on the tissue is randomly adjusted. This is done by decomposing the RGB color of the tissue into H&E color space [33], followed by multiplying the magnitude of H and E of every pixel by two i.i.d. Gaussian random variables with expectation equal to one.

5.1.3 CNN architecture

Following previous work [21], the architecture of our CNN is shown in Tab. 1. We used the CAFFE tool box [18] for the CNN implementation. The network was trained on a single NVidia Tesla 40K GPU.

5.1.4 Experiment setup

To test our method, WSIs of 80% of the patients are randomly selected to train the model and the remaining 20% for testing. Depending on the method, the training set is further divided into two parts randomly: the CNN training set and the logistic regression cross-validation set. The data separation is performed twice and the results are averaged.

The algorithms tested are listed below.

1. NM-LBP: Nuclei Morphological features [12] and rotation invariant Local Binary Patterns [27] are extracted from all the patches extracted from the image, followed by SVM with radial basis function kernel [7]. We report this as a non-CNN baseline.
2. CNN-Vote: CNN followed by voting. Instead of training the second-level logistic regression, the final predicted label of the WSI is determined by voting from the prediction of all its patches.
3. CNN-SMI: CNN followed by max-pooling. This method follows the SMI assumption. The final predicted label of the WSI equals to the predicted label of the patch with maximum probability over all other patches and classes.
4. CNN-LR: CNN followed by our second-level multi-class logistic regression to predict the image-level label. One tenth of the images is held out from the CNN to train the second-level multi-class logistic regression by 10-fold cross-validation.
5. EM-CNN-Vote: CNN-Vote with our EM-based MIL. The percentile for class-level threshold $P_1 = 5\%$. The percentile for image-level threshold $P_2 = 30\%$. See Sec. 3 for the description of these parameters. In each M step, the CNN is trained on all the discriminative patches for 2 epochs.




Figure 2: Some 20X sample patches of 6 types of gliomas from the TCGA dataset. The two patches in each column belong to the same class. Notice the large intra-class heterogeneity.

6. EM-CNN-SMI: CNN-SMI with our EM-based MIL.
7. EM-CNN-LR: CNN-LR with our EM-based MIL.
8. NS-CNN-LR: No Gaussian Smoothing is applied to estimate $P(H | X)$. Other parts are the same as EM-CNN-LR.

Algorithms	Accuracy	mAP	Training Hrs
Chance	0.513	0.689	-
NS-LBP	0.629	0.734	32
CNN-Vote	0.710	0.812	210
CNN-SMI	0.710	0.822	210
CNN-LR	0.752	0.847	201
EM-CNN-Vote	0.733	0.837	243
EM-CNN-SMI	0.719	0.823	243
EM-CNN-LR	0.771	0.845	240
NS-CNN-LR	0.745	0.832	225

Table 3: Glioma classification results. All algorithms in this table use multiscale patches (5X and 20X). The proposed method EM-CNN-LR outperforms CNN-Vote and CNN-SMI significantly.

5.1.5 Results

The results of our experiments are shown in Tab. 3. The confusion matrix is given in Tab. 4. From the confusion matrix, we observe that classification accuracy between GBM and LGG is 97% (chance was 51.3%). A fully supervised method achieved 85% accuracy using a domain specific algorithm trained on ten manually labeled patches per class [25]. To the best of our knowledge our method is the

Ground Truth	Predictions					
	GBM	OD	OA	DA	AA	AO
GBM	214		2		1	
OD	1	47	22	2		1
OA	1	18	40	8	3	1
DA	3	9	6	20		1
AA	3	2	3	3	4	
AO	2	2	3			1

Table 4: The confusion matrix of glioma classification. The Oligoastrocytoma introduces the most confusions due to its nature. See Sec. 5.1.5 for details.

first to classify automatically LGG subtypes, a much more challenging classification task than the benchmark GBM vs. LGG classification. We achieve 57.1% LGG-subtype classification accuracy with chance being 36.7%. Notice that there are more confusions related to Oligoastrocytoma (OA) because it is a mixed glioma that is challenging even for pathologists to agree on. An experiment showed that the inter-observer agreement of four experienced pathologists in a LGG-labeling task similar to the one we are attempting in this paper¹ was approximately 52% and that even after reviewing the cases together, they agreed only around 69% of the time [17].

We tried different patch sizes S and report training time as well and accuracy in Fig. 4. Examples of discriminative and non-discriminative patches identified by the E-step are shown Fig. 3.

¹ Our results are not directly comparable due to different composition of datasets.




Figure 3: Examples of automatically identified discriminative (first row) and non-discriminative (second row) patches of four glioma types.




Figure 4: Training time and results for different sizes of patches. Note that the multiscale settings and EM-based approaches increase accuracy significantly.

5.2. Two-person interaction classification

We performed an additional experiment on the SBU Kinect interaction dataset [40], a less challenging, well understood dataset, both as a sanity check and to test the applicability of our method to other types of large scale, hard to annotate precisely image-based data. The goal of the dataset is to detect two-person interaction activities in RGB-D videos. We use the noisy version of the dataset [40]. Compared to the clean version, each video segment contains five more frames at both its start and end. In total 260 video segments of eight two-person interaction types are provided: “approaching”, “departing”, “kicking”, “punching”, “pushing”, “hugging”, “shaking hands”, and “exchanging”. The classification task is an MIL problem because some of the frames do not contain the labeled video-level actions.

To train the CNN model, we only use depth frames. The architecture of our CNN is similar to that in Table 1 but has much smaller filters and fewer feature maps. The training error converges in a few hours on a single NVidia Tesla

40K GPU. The accuracy of the interaction classification is shown in Table 5. An analysis of the non-discriminative frame identification results is given in Figure 5. Since this is an “easy” dataset, CNN performance is already high and so our EM-based method does not improve them further in this dataset. Nevertheless, it identifies non-discriminative patches effectively. Since there is no evidence showing that non-discriminative frames would contribute to the video-level label in this dataset, there is no need to apply a second-level logistic regression. In contrast, for the glioma classification task, a WSI of OA may contain two non-overlapping regions of different gliomas (See App. A); therefore, patches from these two regions decide the image-level label jointly, and the second level LR is essential.

Algorithms	Accuracy
Chance	0.165
MILBoost	0.873
MILBoost clean	0.911
SVM Linear	0.687
SVM Linear clean	0.876
CNN-Vote	0.926
CNN-SMI	0.920
CNN-LR	0.803
EM-CNN-Vote	0.923
EM-CNN-SMI	0.931
EM-CNN-LR	0.824

Table 5: The results of two-person interaction classification. The SVM and MILBoost algorithm [41] uses skeleton features extracted from the RGB-D frames [40]. Notice that CNNs on the noisy dataset already outperform MILBoost on the clean dataset (MILBoost clean). See text for details.

6. Conclusions

We proposed a Multiple Instance Learning (MIL) based Convolutional Neural Network (CNN) model with a second-level logistic regression that can be applied to high resolution image classification. In this method, we assume that images are too large to apply CNN on whole images. On the other hand, patches extracted from an image may not be discriminative. Thus we follow a Multiple Instance Learning paradigm. We model whether a patch is discriminative or not by a hidden variable and apply EM to maximize the data likelihood. There are two key ideas of our algorithm. First, a spatial prior regularizes the hidden variables. Second, the second-level model is applied on datasets that do not meet the Standard Multiple Instance (SMI) assumption. With our algorithm, we can classify three major subtypes of Low-Grade Glioma (LGG) with an accuracy similar to inter-observer agreement between experienced pathologists. As part of future work we plan to lever-




Figure 5: The localization of identified non-discriminative (irrelevant) frames. The identified non-discriminative frames which do not contain the labeled video-level action are concentrated at either the start or end of the video. This phenomenon is clearer if temporal Gaussian smoothing is applied (EM-CNN-SMI) compared to the algorithm without smoothing (NS-CNN-SMI). In this noisy version of the two-person interaction dataset, each video segment contains five more frames at both its start and end. Therefore, the non-discriminative frames are detected at the two ends of the video as expected, which acts as a sanity check for our claim of detecting non-discriminative frames.

age the non-discriminative patches as part of the dataset. We will explore soft assignments of for the probability that a patch is discriminative. We also plan to test our method with images from other cancer types.

Appendices

A. Description of gliomas in the TCGA dataset

The Cancer Genome Atlas (TCGA) project has whole slide tissue images from glioblastoma multiforme (GBM) and lower grade glioma (LGG) patients. Gliomas are differentiated into grades and sub-categories based on pathologic evaluation of tissues. This evaluation includes assessment of cell characteristics (such as shape and texture) as well as tissue region characteristics (such as existence of necrotic regions). Below is the list of glioma grades and subtypes available at the TCGA repository.

1. GBM: Glioblastoma, ICD-O 9440/3, WHO grade IV. GBM is the high grade glioma that are mostly found in cerebral hemispheres. A Whole Slide Image (WSI) is classified as GBM if and only if one patch can be classified as GBM with high confidence by a pathologist.
2. OD: Oligodendroglioma, ICD-O 9450/3, WHO grade

II.

3. OA: Oligoastrocytoma, ICD-O 9382/3, WHO grade II; and Anaplastic oligoastrocytoma, ICD-O 9382/3, WHO grade III. These two subtypes are mixed in one class because they have the same International Classification of Diseases for Oncology (ICD-O) code. Also, the two subtypes are both mixed gliomas of OD and DA. In other words, a WSI of OA may contain two distinct regions of OD and DA. OA is very hard to classify even by pathologists [17].
4. DA: Diffuse astrocytoma, ICD-O 9400/3, WHO grade II.
5. AA: Anaplastic astrocytoma, ICD-O 9401/3, WHO grade III.
6. AO: Anaplastic oligodendroglioma, ICD-O 9451/3, WHO grade III.

References

- [1] The cancer genome atlas. <https://tcga-data.nci.nih.gov/tcga/>.
- [2] D. Altunbay, C. Cigir, C. Sokmensuer, and C. Gunduz-Demir. Color graphs for automated cancer diagnosis and grading. *Biomedical Engineering, IEEE Transactions on*, 57(3):665–674, 2010.
- [3] J. Amores. Multiple instance classification: Review, taxonomy and comparative study. *Artificial Intelligence*, 201:81–105, 2013.
- [4] S. Andrews, I. Tsochantaridis, and T. Hofmann. Support vector machines for multiple-instance learning. In *NIPS*, pages 561–568, 2002.
- [5] Y. Bengio, A. Courville, and P. Vincent. Representation learning: A review and new perspectives. *PAMI*, 35(8):1798–1828, 2013.
- [6] C. M. Bishop et al. *Pattern recognition and machine learning*, volume 4. 2006.
- [7] C.-C. Chang and C.-J. Lin. Libsvm: a library for support vector machines. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2(3):27, 2011.
- [8] H. Chang, Y. Zhou, A. Borowsky, K. Barner, P. Spellman, and B. Parvin. Stacked predictive sparse decomposition for classification of histology sections. *International Journal of Computer Vision*, pages 1–16, 2014.
- [9] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Semantic image segmentation with deep convolutional nets and fully connected crfs. *arXiv preprint arXiv:1412.7062*, 2014.
- [10] Z. Chen, Z. Chi, H. Fu, and D. Feng. Multi-instance multi-label image classification: A neural approach. *Neurocomputing*, 99:298–306, 2013.
- [11] D. C. Cireřan, A. Giusti, L. M. Gambardella, and J. Schmidhuber. Mitosis detection in breast cancer histology images with deep neural networks. In *MICCAI*, pages 411–418. 2013.

- [12] L. A. Cooper, J. Kong, D. A. Gutman, F. Wang, J. Gao, C. Appin, S. Cholleti, T. Pan, A. Sharma, L. Scarpace, et al. Integrated morphologic analysis for the identification and characterization of disease subtypes. *Journal of the American Medical Informatics Association*, 19(2):317–323, 2012.
- [13] E. Cosatto, P.-F. Laquerre, C. Malon, H.-P. Graf, A. Saito, T. Kiyuna, A. Marugame, and K. Kamijo. Automated gastric cancer diagnosis on h&e-stained sections; ltraining a classifier on a large scale with multiple instance machine learning. In *SPIE Medical Imaging*, pages 867605–867605, 2013.
- [14] A. Cruz-Roa, A. Basavanahally, F. González, H. Gilmore, M. Feldman, S. Ganesan, N. Shih, J. Tomaszewski, and A. Madabhushi. Automatic detection of invasive ductal carcinoma in whole slide images with convolutional neural networks. In *SPIE Medical Imaging*, pages 904103–904103, 2014.
- [15] T. G. Dietterich, R. H. Lathrop, and T. Lozano-Pérez. Solving the multiple instance problem with axis-parallel rectangles. *Artificial intelligence*, 89(1):31–71, 1997.
- [16] J. Foulds and E. Frank. A review of multi-instance learning assumptions. *The Knowledge Engineering Review*, 25(01):1–25, 2010.
- [17] M. Gupta, A. Djalilvand, and D. J. Brat. Clarifying the diffuse gliomas an update on the morphologic features and markers that discriminate oligodendroglioma from astrocytoma. *American journal of clinical pathology*, 124(5):755–768, 2005.
- [18] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014.
- [19] M. Kandemir, C. Zhang, and F. A. Hamprecht. Empowering multiple instance histopathology cancer diagnosis by cell graphs. In *MICCAI*, pages 228–235, 2014.
- [20] M. Kim and F. Torre. Gaussian processes multiple instance learning. In *ICML*, pages 535–542, 2010.
- [21] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, pages 1097–1105, 2012.
- [22] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [23] C. H. Li, I. Gondra, and L. Liu. An efficient parallel neural network-based multi-instance learning algorithm. *The Journal of Supercomputing*, 62(2):724–740, 2012.
- [24] O. Maron and T. Lozano-Pérez. A framework for multiple-instance learning. *NIPS*, pages 570–576, 1998.
- [25] H. S. Mousavi, V. Monga, G. Rao, and A. U. Rao. Automated discrimination of lower and higher grade gliomas based on histopathological image analysis. *Journal of pathology informatics*, 6, 2015.
- [26] M. H. Nguyen, L. Torresani, F. De La Torre, and C. Rother. Weakly supervised discriminative localization and classification: a joint learning process. In *ICCV*, pages 1925–1932, 2009.
- [27] T. Ojala, M. Pietikainen, and T. Maenpaa. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *PAMI*, 24(7):971–987, 2002.
- [28] M. Oquab, L. Bottou, I. Laptev, and J. Sivic. Weakly supervised object recognition with convolutional neural networks.
- [29] G. Papandreou, L.-C. Chen, K. Murphy, and A. L. Yuille. Weakly-and semi-supervised learning of a dcnn for semantic image segmentation. *arXiv preprint arXiv:1502.02734*, 2015.
- [30] D. Pathak, E. Shelhamer, J. Long, and T. Darrell. Fully convolutional multi-class multiple instance learning. *arXiv preprint arXiv:1412.7144*, 2014.
- [31] P. O. Pinheiro and R. Collobert. Weakly supervised semantic segmentation with convolutional networks. *arXiv preprint arXiv:1411.6228*, 2014.
- [32] J. Ramon and L. De Raedt. Multi instance neural networks. 2000.
- [33] A. C. Ruifrok and D. A. Johnston. Quantification of histochemical staining by color deconvolution. *Analytical and quantitative cytology and histology/the International Academy of Cytology [and] American Society of Cytology*, 23(4):291–299, 2001.
- [34] T. H. Vu, H. S. Mousavi, V. Monga, U. Rao, and G. Rao. Dfdl: Discriminative feature-oriented dictionary learning for histopathological image classification. *arXiv preprint arXiv:1502.01032*, 2015.
- [35] N. Weidmann, E. Frank, and B. Pfahringer. A two-level learning method for generalized multi-instance problems. In *Machine Learning: ECML 2003*, pages 468–479. 2003.
- [36] D. H. Wolpert. Stacked generalization. *Neural networks*, 5(2):241–259, 1992.
- [37] Y. Xu, T. Mo, Q. Feng, P. Zhong, M. Lai, E. I. Chang, et al. Deep learning of feature representation with multiple instance learning for medical image analysis. In *ICASSP*, pages 1626–1630, 2014.
- [38] Y. Xu, J.-Y. Zhu, I. Eric, C. Chang, M. Lai, and Z. Tu. Weakly supervised histopathology cancer image segmentation and classification. *Medical image analysis*, 18(3):591–604, 2014.
- [39] K. Yun, J. Honorio, D. Chattopadhyay, T. L. Berg, and D. Samaras. Two-person interaction detection using body-pose features and multiple instance learning. In *CVPR Workshop*, 2012.
- [40] K. Yun, J. Honorio, D. Chattopadhyay, T. L. Berg, and D. Samaras. Two-person interaction detection using body-pose features and multiple instance learning. In *CVPR Workshop*, pages 28–35, 2012.
- [41] C. Zhang, J. C. Platt, and P. A. Viola. Multiple instance boosting for object detection. In *NIPS*, pages 1417–1424, 2005.
- [42] Q. Zhang and S. A. Goldman. Em-dd: An improved multiple-instance learning technique. In *NIPS*, pages 1073–1080, 2001.
- [43] Y. Zhou, H. Chang, K. Barner, P. Spellman, and B. Parvin. Classification of histology sections via multispectral convolutional sparse coding. In *CVPR*, pages 3081–3088, 2014.
- [44] Z.-H. Zhou and M.-L. Zhang. Neural networks for multi-instance learning. In *Proceedings of the International Conference on Intelligent Information Technology*, pages 455–459, 2002.