

In Defense of LSTMs for addressing Multiple Instance Learning Problems

Kaili Wang¹, Jose Oramas², and Tinne Tuytelaars¹

¹ KU Leuven, ESAT-PSI

{first name.last name}@esat.kuleuven.be

² University of Antwerp, imec-IDLab

{first name.last name}@uantwerpen.be

Abstract. LSTMs have a proven track record in analyzing sequential data. But what about unordered instance bags, as found under a Multiple Instance Learning (MIL) setting? While not often used for this, we show LSTMs excel under this setting too. In addition, we show that LSTMs are capable of indirectly capturing instance-level information using only bag-level annotations. Thus, they can be used to learn instance-level models in a weakly supervised manner. Our empirical evaluation on both simplified (MNIST) and realistic (Lookbook and Histopathology) datasets shows that LSTMs are competitive with or even surpass state-of-the-art methods specially designed for handling specific MIL problems. Moreover, we show that their performance on instance-level prediction is close to that of fully-supervised methods.

1 Introduction

Traditional single-instance classification methods focus on learning a mapping between a feature vector (extracted from a single instance) w.r.t. a specific class label. In a complementary fashion, Multiple Instance Learning (MIL) [1] algorithms are tasked with learning how to associate a set of elements, usually referred to as a "bag", with a specific label. In comparison, MIL methods usually require weaker supervision in the form of bag-level labels. The MIL problem has a long history, and various solutions have been proposed over time. Here, we advocate the use of standard LSTM networks in this context, as a strong baseline, yielding competitive results under a wide range of MIL settings.

Long short-term memory (LSTM) networks [2] have been proposed as an extension over standard recurrent neural networks, to store information over long time intervals in sequential data. They have been used extensively and very successfully for modeling sentences (sequences of words) in text documents [3], e.g. for machine translation [4] or sentiment analysis [5]. Later, they have been employed in several other fields including computer vision [6–10] and speech processing [11, 3]. LSTMs provide great flexibility for handling data sequences. There is no need to know the length of the sequence beforehand, and they can operate on sequences of variable size. In addition, they are capable of accumulating information by using a memory mechanism with add/forget[12] functionality.

As they need the input to be provided as a sequence, LSTMs do not seem an appropriate choice for analyzing unordered bags at first - but is that so ? Obviously, the capability to remember the temporal (order) information can be attributed to the *memory* ability of LSTMs. This memory ability is capable of capturing other types of information beyond order as well. Take the LSTMs used for action recognition as an example. For some finegrained actions (e.g. opening or closing a door), the order of the events is key and this is picked up by the LSTM. However, for other actions the context information provides the most important cue (e.g. playing tennis or cooking). This does not depend on the temporal order, but can still be learned using LSTM.

Starting from an unordered bag, we can always transform it into a sequence by imposing a random order on the instances, making it suitable for LSTMs. The order is not relevant, but that does not matter: the LSTM can still process the data and extract useful information from it. In fact, this is also how humans often deal with unordered bags: e.g. if one is asked to count the number of rotten apples in a basket, most of us would just pick the apples one by one, in random order, inspect them and keep track of the count. The order does not matter, but nevertheless, treating them in a sequential order comes very naturally.

The observations above clearly hint to a promising capability of LSTMs for addressing MIL problems. Yet, LSTMs are not often used in this way (see our related work section for a few notable exceptions). Therefore, we present a systematic analysis on the performance of LSTMs when addressing MIL problems. More specifically, we conduct a series of experiments considering different factors that may affect the performance of LSTMs. First, we consider the standard MIL problem [13–15], with bags of instances without sequential order. Second, we study the effect of the order in which the instances of each bag are fed to the LSTM network. Likewise, in a third test, we investigate the influence of the cardinality (size) of the bag on performance. Fourth, we assess the effect that the complexity of the data has on the previous observations. Toward this goal we conduct experiments considering bags derived from the MNIST dataset [16], clothing-item images from the Lookbook dataset [17] and Histopathology images [18]. Fifth, we inspect how the internal state of the LSTM changes when observing each of the instances of the bag. Moreover, we propose an LSTM-based framework that can predict the instance-level labels by only using bag-level labels, in a weakly supervised manner.

Our contributions are three-fold: i) We advocate the application of LSTM networks on general MIL problems, as a means for encoding more general underlying structures (i.e. not limited to ordered data) within bags of instances. ii) We conduct an extensive systematic evaluation showing that LSTMs are capable of capturing information within bags that go beyond ordered sequences. Moreover, we show that their performance is, surprisingly, comparable or even superior to that of methods especially designed to handle MIL problems. iii) We propose a framework for weakly supervised learning based on LSTM, capable of modeling distributions at the instance-level, using only bag-level annotations.

2 Related Work

Efforts based on LSTMs/RNNs aiming at modelling unordered bags are quite rare. [19, 20] propose a memory-based recurrent architecture and apply it on sorting [19] and traversal of graphs [20], where the input data can be regarded as unordered bags. [21] considers a fashion outfit to be a sequence (from top to bottom) and each item in the outfit as an instance. Then, a LSTM model is trained to sequentially predict the next item conditioned on previous ones to learn their compatibility relationships. Later, [22] used a CNN-RNN model to perform multiple label predictions, where LSTMs were used to decode the labels in an unordered manner. Different from this work which focused on the decoding part, we investigate the encoding of the unordered bags. [23] proposed to learn permutation-invariant bag representations by summing features across instances and applying non-linear transformations. This can be regarded as a specific case of [18] where the weight of the instances are uniform. [24] uses LSTMs to capture the function over the sets, which is different from ours, whereas we use LSTM to model the set representation and learn instance-level distributions from the set label. These works either use LSTMs to handle unordered data on some specific settings [21, 23] or use them just as side experiments [19, 20]. Here, we propose the use of LSTMs to address more *general* MIL problems.

On the task of modeling general bag representations, we position our approach w.r.t. efforts based on neural networks, specifically those with deep architectures since our work is based on LSTMs. Please refer to [13–15] for detailed surveys covering non-deep methods. [25] proposed a multiple instance neural network to estimate instance probabilities. This idea is further extended in [26] which uses a neural network to learn a bag representation and directly carry out bag classification without estimating instance-level probabilities or labels. In parallel, [18] proposed an attention mechanism to learn a pooling operation over instances. The weights learned for the attention mechanism on instances can serve as indicators of the contribution of each instance to the final decision – thus, producing explainable predictions. [27] proposed a similar idea, using the computed bag representations, to measure distances between image bags. [28] proposed to update the contributions of the instances by observing all the instances of the bag a predefined number of iterations. Along a different direction, [29] proposed a hierarchical bag representation in which each bag is internally divided into subbags until reaching the instance level. Very recently, [30] proposed to consider the instances in the bags to be non-i.i.d. and used graph neural networks to learn a bag embedding.

Similar to [18, 26] we embed the instance features from each bag into a common space and the bag representation is used to make direct bag predictions. Similar to [30, 28] we aim at learning the underlying structure within the bags. Different from [30], our method does not rely on hand-tuned parameters, e.g. manual graph construction. Moreover, the improvement in performance displayed by our method is not sensitive to the possible lack of structure within each bag. Compared to [28], our method only requires *a single pass* through all the instances. Moreover, our method is able to go beyond binary classifica-

tion tasks and handle more complex classification and regression tasks. Finally, most of the works mentioned above operate under the standard Multiple Instance (MI) assumption. In contrast, the proposed approach is able to learn the underlying structure of bags of instances, thus, being robust to several MI assumptions/problems [15].

3 Methodology

We begin our analysis by defining the different parts that compose it. First, we formally define MIL problems and draw pointers towards different MI assumptions that they commonly consider. Then, we introduce the LSTM-based pipeline that will be considered to model bags of instances throughout our analysis.

3.1 Underlying Structures within Bags of Instances

As was said earlier, underlying sequential structures between the instances within a bag is a cue that LSTMs are capable of encoding quite effectively. In fact, this capability have made them effective at handling problems defined by these sequences, e.g. actions, speech, text. However, this sequential order is just one of many possible underlying structures that could be present between the instances within the bags processed by a LSTM. Encoding these underlying structures and making predictions about bags of instances is the main objective of Multiple Instance Learning (MIL) [13–15]. We will conduct our analysis from the perspective of MIL, where LSTMs will play an active role in modeling the underlying bag structure.

Given the bag $X_j = \{x_1, x_2, \dots, x_m\}$ of instances x_i with latent instance-level labels $C_j = \{c_1, c_2, \dots, c_m\}$, traditional MIL problems aim at the prediction of bag-level labels y_j for each bag X_j . The MIL literature covers several underlying bag structures, referred to as *assumptions*, that have been commonly considered in order to define bag-level labels. We refer the reader to [13–15] for different surveys that have grouped these assumptions based on different criteria.

3.2 Proposed Pipeline

The proposed pipeline consists of three main components. Given a bag X of m instances x_i , each of the instances x_i is encoded into a feature representation f_i through the *Instance Description Unit*. Then, each element is fed to the *Iterative bag Pooling Unit*, producing the aggregated bag representation S . Finally, a prediction \hat{y} is obtained by evaluating the bag representation via the *Prediction Unit*.

Instance Description Unit This component receives the bag of instances in raw form, i.e. each of the instances $x_i \in \mathbb{R}^{[d]}$ that compose it, in its original format. It is tasked with encoding the input bag data into a format that can be

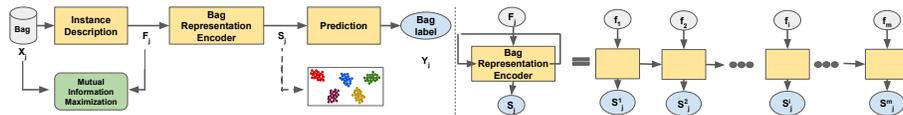


Fig. 1: Left: Proposed approach pipeline. Right: Iterative bag pooling unit. The bag representation S_j^i is updated each time the representation f_i of an element is observed.

processed by the rest of the pipeline. As such, it provides the proposed method with robustness to different data formats/modalities. More formally, given a dataset $\{X_j, y_j\}$ of bags X_j paired with their corresponding bag-level labels y_j , each of the bags X_j is encoded into a feature $F_j = \{f_1, f_2, \dots, f_m\}$. This is achieved by pushing each of the instances x_i that compose it, through a feature encoder $\tau(\cdot)$ producing the instance-level representation $f_i = \tau(x_i), f_i \in \mathbb{R}^{[n]}$. Selection of this component depends on the modality of the data to be processed, e.g. VGG [31] or ResNet [32] features for still images, Word2Vec [33] or BERT [34] for text data, or rank-pooled features [35] or dynamic images [36] for video data.

Maximizing Mutual Information from Instances Mutual information can be used to measure the (possibly non-linear) dependency between two variables, noted as $I(A; B)$. Maximizing mutual information between input instance and its representation helps the model learn a better representation [37][38]. It is useful as a regularizer especially when learning a model from scratch. In our method, we follow [37] where the total objective function is:

$$L = \alpha \cdot \max(MI_{global}) + \beta \cdot \max(MI_{local}) + \gamma \cdot \text{PriorMatching} \quad (1)$$

MI_{global} and MI_{local} are the global and local Mutual information where the latter one is calculated between intermediate feature map and final representation. *PriorMatching* is used to match the output of the IDU to a prior: combined with the maximization of mutual information, Eq. 1 can constrain representations according to desired statistical properties. For the sake of space, please refer to [37] and the supplementary material for more details regarding the derivation.

Bag Representation Encoder The main goal of this component is to derive a bag-level representation $S_j \in \mathbb{R}^{[n]}$ that is able to encode all the instances x_i , and any possible underlying structure between them. As mentioned earlier, LSTM is utilized to model the underlying [unordered] structure in the bag.

We aim at learning a bag representation that is independent of both the cardinality m of the bag and the nature of the underlying structure. Starting from the element-level representations F_j computed in the previous step, this is achieved by iteratively looking at the representations f_i , from each of the instances x_i , one at a time. In each iteration i an updated bag-level representation S_j^i is computed. In parallel, following the LSTM formulation, a feedback loop provides information regarding the state of the bag representation that will be

considered at the next iteration $i+1$. Finally, after observing all the m instances x_i in the bag, the final bag representation $S_j^m|_{i=m}$ is taken as the output S_j of this component.

The notion behind this iterative bag pooling idea is that instances observed at specific iterations can be used to compute a more-informed bag-level representation at later iterations. Thus, allowing to encode underlying relationships or structures among the instances of the bag. While this iterative assumption may hint at a sequence structure requirement within each bag, our empirical evaluation strongly suggests this not to be the case. (see Sec. 4.2)

In practice, we use Bi-directional LSTMs which observe the instances in a bag from the left-to-right and right-to-left directions. This will further ensure that the context in which the instances of the bag occur is properly modelled.

Prediction Unit Having a bag-level representation S_j for bag X_j , this component is tasked with making a bag-level prediction $\hat{y}_j=g(S_j)$ that will serve as final output for the pipeline. The selection of the prediction function $g(\cdot)$ is related to the task of interest. This unit provides our method with flexibility to address both classification and regression tasks.

4 Analysis

4.1 What kind of information can be captured by LSTMs?

This experiment focuses on performing multiple instance predictions based on visual data. Following the protocol from [18] we use images from the MNIST dataset [16] to construct image bags to define four scenarios, each following a different assumption: *Single digit occurrence*, *Multiple digit occurrence*, *Single digit counting* and *Outlier detection*.

For this series of experiments, we use a LeNet as Instance Descriptor unit (IDU) and an LSTM with an input and cell state with 500 dimensions as Bag Representation Encoder (BRE), respectively. Both the IDU and BRE components are trained jointly from scratch. We compare the obtained performance w.r.t. the attention-based model from [18] and the dynamic pooling method from [28]. Mean error rate in the binary classification task is adopted as performance metric in these experiments. Please note, we *do not* traverse all the possible permutations in the following experiments. On the contrary, only a small proportion of them are seen by the model.

The main objective of this experiment is to answer the following questions: i) whether other underlying bag structures, outside of sequential order, can be encoded properly by LSTMs?, and ii) *how competitive are LSTMs when compared with methods from the MIL literature specifically designed for modeling the underlying bag structures?*

Single Digit Occurrence In this scenario we follow the standard MI assumption and label a bag as positive if at least one digit '9' occurs in the bag. The

digit '9' is selected since it can be easily mistaken with digit '4' and '7' [18], thus, introducing some instance-level ambiguity. We define bags with mean cardinality $m=10$, and verify the effect that m has on performance by testing two standard deviation values, $\sigma=2$ and $\sigma=8$. We repeat this experiment five times generating different bags and weight initializations. We report mean performance in Table 1 (col. II and III).

Discussion: The results indicate that, in this task, our performance is comparable with the state-of-the-art for lower values of σ and superior as σ increases. This is to some extent expected, since at lower σ the cardinality (i.e. the number of instances) of each bag is almost fixed. This setting is favorable for the attention-based method since it operates in a feed-forward fashion. Yet, note the high standard deviation in performance produced by this baseline. On the contrary, at higher σ values there is a higher variation of cardinality across bags. Under this setting, feed-forward approaches start to produce higher errors. Here our method produces superior performance, ~ 1.4 percentage points (pp) w.r.t. to the state-of-the-art.

Method	single digit($\sigma=2$)	single digit($\sigma=8$)	multiple digits	digit counting	outlier detection
Atten. Based	2.8 ± 4.8	4.5 ± 0.4	28.5 ± 0.7	33.4 ± 19.3	37.0*
Gated Atten. Based	4.0 ± 0.9	4.6 ± 0.5	27.4 ± 0.9	11.9 ± 3.6	37.4*
Dyn. Pool	5.6 ± 1.1	6.1 ± 1.2	28.5 ± 6.6	25.4 ± 1.8	40.9*
Ours w/o Mut. Info.	3.5 ± 1.1	3.1 ± 0.5	6.4 ± 1.4	9.0 ± 2.7	50.0
Ours	4.0 ± 0.4	4.1 ± 1.4	3.5 ± 1.3	7.4 ± 1.2	2.07

Table 1: Mean error rate (in percentage points) of experiments considering digits from the MNIST dataset. (*) refers to baselines which include the Mutual Information loss.

Multiple Digit Occurrence This is an extension of the previous scenario in which instead of focusing on the occurrence of a single digit class, the model should recognize the occurrence of instances of two digit classes (presence-based MI assumption [15]). More specifically, a bag is labeled positive if both digits '3' and '6' occur in it, without considering the order of occurrence. For this scenario 1,000 bags are sampled for training. Results are reported in Table 1 (col. IV).

Discussion: It is remarkable that when making this simple extension of considering the occurrence of multiple digits, i.e. '3' and '6', the state-of-the-art methods suffer a significant drop in performance. This drop put the state-of-the-art methods ~ 27 pp below, on average, w.r.t. the performance of our method. Please note that in this experiment the order (or location) of the two digits does *not* matter. This supports previous observations that LSTMs can indeed handle multiple instances of interest, independent of the ordering in which they occur within the bags. In this scenario, where observing multiple instances is of interest, the model needs to “remember” the information that it has seen in order to assess whether instances of the classes of interest have been encountered. The feed-forward models lack information persistence mechanisms; which translates

to a poor ability to remember and to handle multiple instances of interest. Surprisingly, in spite of its iterative nature, the Dynamic pooling method is not able to preserve the information it has observed across iterations, resulting in similar performance as the other baselines.

Digit Counting Previous scenarios addressed the classification task of predicting positive/negative bag-level labels. In contrast, in this scenario, we focus on the regression task of counting the number of instances of a specific digit class of interest within the bag (presence-based MI assumption). In order to make our approach suitable to address a regression problem, instead of using a classifier as prediction unit we use a regressor whose continuous output is rounded in order to provide a discrete count value as output. In this experiment the digit '9' is selected as the class to be counted. The mean cardinality of each bag is fixed to $m=15$. Performance is reported in Table 1 (col. V).

Discussion: From Table 1 (col. V) the same trend can be observed: our method has superior performance and higher stability than the attention-based model and other baselines. When conducting this counting task, our method obtains a performance that is superior by 24 pp w.r.t. the attention-based model and by 16 pp w.r.t. the dynamic pooling. These results further confirm the capability of LSTMs to handle this type of unordered regression problems [39].

Digit outlier detection This task is concerned with identifying whether a bag contains a digit which is different from the majority (outlier). Different from *Single digit occurrence*, this task is more difficult since the model has to understand: i) the two digit classes that might be present in the bag, and ii) the proportion condition that makes the bag an outlier. This is different from *Single digit occurrence* where it only needed to identify the "witness" digit '9'. Besides, there is no restriction on the outlier and majority digits, they can be any digit class from MNIST dataset. This constitutes a collective MI assumption since all the instances determine the underlying structure of the bag. Therefore, given the complexity of this task, in this experiment we apply the mutual information loss on every baseline method in order to assist their training. We use 10,000 bags to train the model and 2,000 bags to test. The bag cardinality is 6 with 1 standard deviation. Table 1 (col. VI) shows quantitative results of this experiment.

Discussion: It is remarkable that, even after applying the mutual information loss on the other baselines, they still have a low performance on this task. We notice that the Attention and Gated Attention methods work slightly better than Dynamic Pooling. More importantly, our method, based on LSTMs, outperforms the baselines by a large margin (~ 36 pp). This suggests that LSTMs are quite capable at modeling this type of bag structure, even to the point of outperforming MIL methods tailored to model bag-based structures.

4.2 Does the result depend on the order chosen?

The short answer is no. The reason is that in the training phase we push the bags with different orders (as a form of data augmentation) to the model while

the bag labels are the same. By following this procedure the loss function will not penalize differences in the order in which the instances are observed. To further verify this, we repeated the test phase of our experiments 100 times with the contents of each bag (cardinality= m , in total $m!$ combinations) shuffled thus producing bags with 100 different orders. Then, similar to Sec. 4.1, we measure the error rate and report the mean performance. The obtained error rate is $(4.2\pm 0.6)\%$, $(3.5\pm 0.4)\%$, $(7.8\pm 0.7)\%$ for the Single-digit, Multiple-digit and Digit counting experiments, respectively ($m=10, 12, 15$ respectively). They are very close to the numbers reported in Table 1. The results verify that the LSTM is able to learn, to a good extent, that the underlying MIL assumptions were permutation invariant - changing the order of instances of a bag has a relatively low effect on the prediction in most of the cases.

4.3 Does the result depend on the cardinality of the bag?

No, modeling bag representations via LSTMs seems robust enough to bags with different cardinality (sizes). We verify this by conducting an extended experiment based on the *multiple instance occurrence* scenario. Firstly, we consider bags with higher cardinality but keep *only one relevant instance pair* ('3', '6') present using one of our trained models (mean bag cardinality $m=12$). We obtained error rates of 7%, 14.5%, 42%, and 44% for bag cardinality 20, 50, 100 and 200, respectively. This result is not surprising since during training the bag cardinality was much lower. To have a fair experiment, we use bags with higher cardinality to finetune our model, using 1/5 amount of the original number of training bags (i.e. now we use 200 bags). Similarly, the larger bags still contain only one pair of relevant digits. This results in error rates of $(2.38\pm 0.41)\%$, $(3.13\pm 0.89)\%$, $(4.25\pm 1.3)\%$ for mean bag cardinality of 50, 100 and 200, respectively. This shows that LSTMs are still capable of modeling unordered bags even when bags with significantly higher cardinality are considered, although, unsurprisingly, training and testing conditions should match.

4.4 Effect of the Complexity of the Data

In this section, we shift our analysis to real-world data. Summarizing, the results show that our method still works comparable and even better than the baselines.



Fig. 2: Examples of instances for the original (left), occluded (middle) and database images (right) in our cross-domain clothing retrieval experiment.

Cross-domain clothing retrieval For this experiment, we divide images from the Lookbook dataset into two domains: catalog clothing images and their corresponding human model images where a person is wearing the clothing product, see Fig. 2. Each clothing product has one catalog image and several human model images. We only consider the products with five or more human model images, resulting in 6616 unique products (latent classes c_i) with around 63k images in total. Every product image has 5-55 human model images. The training bag contains 4000 classes while the validation and test bags have 616 and 2000 classes, respectively. We run two experiments on this dataset as described in the following sections. Given the higher complexity of images in this dataset, we use a pre-trained VGG16 [31] as IDU. Since this unit is pretrained, the mutual information loss is not applied for this unit in this experiment. Moreover, we set the dimensionality of the input and cell state of our LSTM to $n=2048$.

For this experiment, human model images are used as queries while catalog images serve as database, thus, defining a many-to-one retrieval setting. The cardinality of each bag is the same as the number of human model images of each product (class). We conduct two variants of this experiment. On the first variant we use the complete image, as it is originally provided. The second is an occluded variant where every human model image in a bag is divided into a 4×4 grid of 16 blocks. 12 of these blocks are occluded by setting all the pixels therein to black. By doing so, every single image in a bag can only show part of the information while their combination (i.e. the whole bag) represents the complete clothing item. Catalog images in the database are not occluded in this experiment. This experiment can be regarded as an extreme case of standard MI assumption, where all the instances in each bag is positive.

As baselines, in addition to the attention-based model we follow DeepFashion [40], and train a model to perform retrieval by computing the distances by considering single image representations instead of bag-based representations. Following the multiple queries approach from [41], we report performance of three variants of this method: *Single-AVE*, where the distance of each bag is computed as the average of the distances from every image in the bag w.r.t. an item in the database; *Single-MIN*, where the distance of the bag is defined as the minimum distance of an image in the bag w.r.t. an item in the database; and *Single Fea. AVE*, where the distance of the bag is calculated as the distance of a prototype element w.r.t. an item in the database. As prototype element we use the average feature representation \bar{f}_i from the representation f_i of every element in the bag. We refer to these baselines as *Single-image models*.

Method	rec.@1	rec.@10	rec.@20	rec.@50
Atten.	13.75	39.25	49.70	63.60
Dyn. Pool	16.75	47.65	59.45	73.60
Single AVE	20.55	57.05	68.25	81.90
Single MIN	22.60	58.15	69.20	82.50
Single Fea. AVE	20.15	56.25	67.85	81.50
Ours w/o mut. info.	22.95	58.65	68.70	83.00

Table 2: Retrieval on the original Lookbook dataset.

Method	rec.@1	rec.@10	rec.@20	rec.@50
Atten.	3.55	20.6	32.95	53.65
Dyn. Pool	1.95	11.95	29.35	32.55
Single AVE	3.65	23.85	35.06	56.10
Single MIN	5.25	26.05	37.35	55.00
Single Fea. AVE	5.10	25.60	36.95	54.65
Ours w/o mut. info.	9.25	34.75	45.00	61.80

Table 3: Retrieval on the occluded Lookbook dataset.

Discussion: Table 2 shows that in the original setting our method tends to obtain superior recall values in the majority of the cases, with the exception of the case when the closest 20 items (recall@20) are considered. When looking at the occluded variant of the experiment, a quick glance at Table 3 shows that, compared to the original setting, absolute performance values on this setting are much lower. This is to be expected since this is a more challenging scenario where the model needs to learn the information cumulatively by aggregating information from parts of different images. In this occluded setting, our method clearly outperforms all the baselines. This could be attributed to the information persistence component from the LSTMs. This component allows our method to select what to remember and what to ignore from each of the instances that it observes when updating the bag representation used to compute distances. The difference w.r.t. to the *Single-AVE* and *Single-MIN* baselines is quite remarkable given that they require a significant larger number of element-wise distance computations w.r.t. items in the database. This may lead to scalability issues when the dataset size increases, as the computation cost will grow exponentially.

Moreover, in both occluded and non-occluded datasets, we notice that the *Single-image model* baselines have a superior performance w.r.t. the attention-based model and dynamic pooling model. We hypothesize that is because the single-image models can better exploit important features, e.g. discriminative visual patches, since they compute distances directly in an instance-wise fashion. In contrast, it is likely that some of these nuances might get averaged out by the feature aggregation step that is present in the attention-based model.

Colon Cancer Prediction This task consists of predicting the occurrence of Colon cancer from histopathology images. The used Colon cancer dataset contains 100 500×500 H&E images with a total of 22k annotated nuclei. There are four types of nuclei: *epithelial*, *inflammatory*, *fibroblast*, and *miscellaneous*. This experiment focuses on identifying whether colon cancer histopathology images contain a specific type of nuclei. We follow the protocol from [18] and treat every H&E image as a bag composed by instances (patches) of 27×27 pixels centered on detected nuclei. The bag cardinality varies from 6 to 796 depending on the number of nuclei present in the image. Following a standard MI assumption, a bag is considered positive if it contains *epithelial* nuclei since Colon cancer originates from epithelial cells [18][42] This produces a dataset with 51 and 48 positive and negative bag examples, respectively. We extend this dataset via data augmentation as in [18].

We adapt an architecture which is similar to [43] to define the IDU and a 512 dimension input and cell state to define the LSTM (BRE). The whole model is trained from scratch. Following the protocol, only bag-level binary labels are used for training. We conduct experiments considering the same baselines as in previous experiments. We apply five-fold cross validation and report the mean performance and standard deviation. For reference, we also provide the results presented in [18] for the baselines Atten.* and Gated Atten.*. Table 4 shows quantitative results in terms of Accuracy and F1-Score.

Discussion: This experiment, where a bag can have up to 796 instances, serves a good test-bed to assess the performance of the proposed method on bags with high cardinality. From the results in Table 4, we can notice that our method still outperforms all the considered baselines.

Metric	Atten.*	Gated Atten.*	Atten.	Gated Atten.	Dyn. Pool	Ours w/o mut. info.	Ours
Accuracy	90.40±1.10	89.80±2.00	88.79±6.16	86.89±3.93	87.89±2.37	90.89±2.06	92.74±2.41
F1-Score	90.10±1.10	89.30±2.20	88.85±6.35	86.87±6.67	88.18±2.11	90.66±2.80	93.08±1.36

Table 4: Colon cancer experiment results.

5 From Internal States to Instance-level Representations

Previous efforts [2, 44, 7] based on ordered bags have shown that the internal state of the representation within the LSTMs can be used to predict future instances. Here, we have shown that LSTMs can also encode other types of bag-based information internally. This begs the question - *what else can the internal representation in LSTMs reveal in the unordered setting?*. Here we conduct an analysis aiming to answer this question.

5.1 Weakly Supervised Learning of Instance-Level Distributions

We have presented using LSTM to make predictions from a bag-level representation S_j through the use of a prediction function $g(\cdot)$. There is a connection between the MIL task and the distribution of the instance representation. Based on this observation we put forward the following hypotheses:

- *Hypothesis 1: A model trained for a MIL task can learn the underlying distribution over the instances.*

- *Hypothesis 2: A prediction function $g(\cdot)$ trained on the bag representation S can be used to make instance-level predictions if the distribution from S , influenced by the underlying MI assumption, is close to that of F .*

We propose the following approach to recover the underlying instance-level representation and make instance-level predictions. We break down the instance bag $X = \{x_1, x_2, \dots, x_m\}$ into m singleton bags $X_1 = \{x_1\}$, $X_2 = \{x_2\}$, ..., $X_m = \{x_m\}$. The singleton bag $X_j = \{x_j\}$ is sent to the model, passing the IDU and the LSTM. Afterwards, the output S_j of the LSTM from every singleton is collected into a feature matrix \mathbb{S} , $\mathbb{S} \in \mathbb{R}^{[m \times n]}$. Then, k-means clustering algorithm is applied on \mathbb{S} with the number of clusters determined by the corresponding MIL task. We use a similar metric to clustering purity, where we calculate the purity of each cluster first and average them instead of calculating the purity of all samples. By doing this we avoid problems caused by imbalanced data. The clustering performance reflects the ability of modeling the distribution of instances for the model.

5.2 Weakly Supervised Instance-level Learning

In Sec. 5.1 we presented two hypotheses related to the weakly supervised instance-level learning. We will address them in this section.

Modelling Instance-level Representations In Sec. 4.1 and 4.4, we trained both IDU and LSTM from scratch by considering the bag-level labels only. This can be regarded as weakly supervised learning if the goal is to make instance-level predictions. For attention-based methods, we collect the output of IDU and multiply with weight 1, since it is a singleton bag and there is no LSTM. Following this procedure, both methods use the features after their respective units handling the MIL task. We evaluate instances from both testing/training bag for the baseline and our model, respectively. We choose the Gated-Attention model as a baseline since it works best among the attention-based methods in Sec. 4.1. Table 5 reports the clustering performance metric described in Sec. 5.1.

Task	Gated Atten. (test/train)	Ours (test/train)
<i>single digit</i> (2 classes)	98.69/98.92	97.59/97.42
<i>multiple digits</i> (3 classes)	85.92/87.47	97.94/97.06
<i>digit counting</i> (2 classes)	99.22/99.31	99.15/99.23
<i>outlier detection</i> (10 classes)	59.33/57.02	97.96/97.52

Table 5: Instance clustering accuracy from MNIST-bag task models.

Method	TP (test/train)	TN (test/train)	mean Acc (test/train)
Atten.	32.42 / 21.25	98.45 / 99.22	65.43 / 63.60
Ours	73.47 / 70.73	92.39 / 92.28	82.93 / 81.51
Supervised	<i>78.92 / 92.09</i>	<i>91.14 / 98.22</i>	<i>85.03 / 95.16</i>

Table 6: Instance label accuracy for Colon cancer dataset.

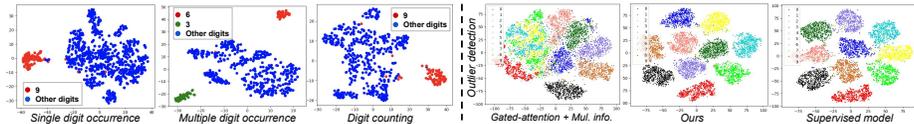


Fig. 3: t-SNE visualization of features extracted from **our** MIL model in three MNIST bag tasks (left) and with baseline models in *Outlier detection* (right).

Discussion: Table 5 indicates that for simple tasks, such as *single digit occurrence* and *digit counting*, both attention-based and our methods can distinguish the background digits and witness digits. To handle the MIL task, the model just needs to differentiate between the witness digit (“9”) from other digits. Therefore, there should only be two clusters/classes. Three clusters/classes are assigned to *multiple digits* because the model needs to distinguish the two witness digits from the others.

For the case of *outlier detection*, in order to detect the outlier(s) from a bag, the model needs to distinguish every digit. For this reason, once capable of handling this MIL task, the models should also have the ability to cluster/classify the 10 digits. It is clear that our model trained for this task has learned very good discriminative features for all 10-class digits, while the attention-based method fails, even when the mutual information loss is still applied on top of it. The clustering accuracy is close to the known performance of ~98% accuracy of the supervised LeNet model [16]. This is strong evidence showing that our method is able to learn an instance-level representation in a weakly supervised manner.

In addition, Fig. 3 shows the t-SNE visualizations for features extracted by our method in the testing bag of the four tasks. The figure shows how discriminative the singleton features are. These results prove that our *Hypothesis 1* is correct.

Instance-level prediction: The colon cancer dataset contains 7,722 epithelial nuclei and 14,721 other nuclei. We select one of the models we trained earlier and treat the patches as singleton bags (i.e. bags only contain one patch). The singleton bags are sent to the model to make instance-level predictions: epithelial or not. In the meantime we also use the same training-test split to train a fully supervised model. We report the instance-level accuracy in Table 6. In addition, Fig. 4 shows the patches that are classified as epithelial nuclei.

Discussion. This task meets the requirement of *Hypothesis 2*: the bag representation S contains the information whether the epithelial nuclei exist in a bag, which is close to what would be expected from instance-level feature F . Our model achieves the best performance for bag-level prediction. It also has a good performance on the instance-level prediction. The mean accuracy is close to the supervised model and significantly better (~ 18 pp) than that of the Attention-based model. It clearly shows that our MIL model can be used to predict the instance labels. In addition, Fig. 4 shows that our model has a better ability to identify the nuclei of interest, which can be useful for pathologists.

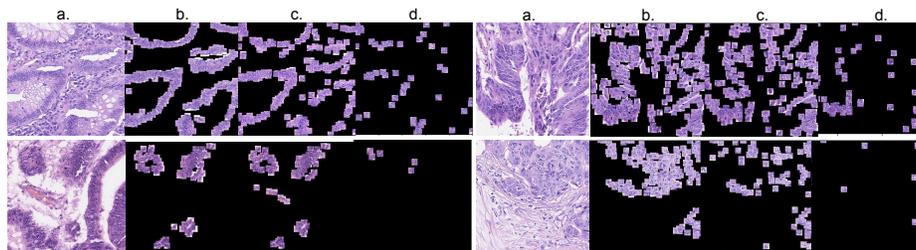


Fig. 4: a) The original H&E image. b) The epithelial nuclei patches (Ground-Truth). c) The epithelial nuclei patches detected by our MIL model. d) The epithelial nuclei patches detected by attention-based MIL model

6 Conclusion

We investigate the potential of LSTMs at solving MIL problems. Through an extensive analysis we have shown that LSTMs can indeed capture additional information of the underlying structure within the bags. Our results suggest that its performance at modeling more general bag structures is comparable and even better than that from methods tailored for MIL problems. Our method can also model the instance-level distribution in a weakly supervised manner.

Acknowledgement This work was partially supported by the KU Leuven AI implants, FWO G.0A47.20N, and a NVIDIA Academic Hardware Grant.

References

1. Sammut, C., Webb, G.I.: Multi-instance learning. *Encyclopedia of Machine Learning* (2011)
2. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Computation* **9** (1997) 1735–1780
3. Sundermeyer, M., Schlüter, R., Ney, H.: Lstm neural networks for language modeling (2012)
4. Sutskever, I., Vinyals, O., Le, Q.V.: Sequence to sequence learning with neural networks. In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N.D., Weinberger, K.Q., eds.: *Advances in Neural Information Processing Systems 27*. Curran Associates, Inc. (2014) 3104–3112
5. Le, Q.V., Mikolov, T.: Distributed representations of sentences and documents. In: *ICML. JMLR Workshop and Conference Proceedings, JMLR.org* (2014)
6. Liu, J., Shahroudy, A., Xu, D., Wang, G.: Spatio-temporal lstm with trust gates for 3d human action recognition. In Leibe, B., Matas, J., Sebe, N., Welling, M., eds.: *ECCV*. (2016)
7. SHI, X., Chen, Z., Wang, H., Yeung, D.Y., Wong, W.k., WOO, W.c.: Convolutional lstm network: A machine learning approach for precipitation nowcasting. In Cortes, C., Lawrence, N.D., Lee, D.D., Sugiyama, M., Garnett, R., eds.: *Advances in Neural Information Processing Systems 28*. Curran Associates, Inc. (2015) 802–810
8. Alahi, A., Goel, K., Ramanathan, V., Robicquet, A., Fei-Fei, L., Savarese, S.: Social lstm: Human trajectory prediction in crowded spaces. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. (2016)
9. Zhu, W., Lan, C., Xing, J., Zeng, W., Li, Y., Shen, L., Xie, X.: Co-occurrence feature learning for skeleton based action recognition using regularized deep LSTM networks. In: *Thirtieth AAAI Conference on Artificial Intelligence*. (2016)
10. Singh, B., Marks, T.K., Jones, M., Tuzel, O., Shao, M.: A multi-stream bi-directional recurrent neural network for fine-grained action detection. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. (2016)
11. Graves, A., Schmidhuber, J.: Framewise phoneme classification with bidirectional lstm and other neural network architectures. *NEURAL NETWORKS* (2005) 5–6
12. Gers, F.A., Schmidhuber, J., Cummins, F.: Learning to forget: Continual prediction with lstm. *Neural Computation* **12** (1999) 2451–2471
13. Amores, J.: Multiple instance classification: Review, taxonomy and comparative study. *Artif. Intell.* **201** (2013) 81–105
14. Carbonneau, M.A., Cheplygina, V., Granger, E., Gagnon, G.: Multiple instance learning: A survey of problem characteristics and applications. *ArXiv abs/1612.03365* (2018)
15. Foulds, J., Frank, E.: A review of multi-instance learning assumptions. *The Knowledge Engineering Review* **25** (2010) 1–25
16. LeCun, Y., Cortes, C.: MNIST handwritten digit database. (2010)
17. Yoo, D., Kim, N., Park, S., Paek, A.S., Kweon, I.: Pixel-level domain transfer. *ECCV* (2016)
18. Ilse, M., Tomczak, J.M., Welling, M.: Attention-based deep multiple instance learning. *arXiv preprint arXiv:1802.04712* (2018)
19. Graves, A., Wayne, G., Danihelka, I.: Neural turing machines. *CoRR abs/1410.5401* (2014)
20. Graves, A., Wayne, G., Reynolds, M., Harley, T., Danihelka, I., Grabska-Barwińska, A., Colmenarejo, S.G., Grefenstette, E., Ramalho, T., Agapiou, J.,

- Badia, A.P., Hermann, K.M., Zwols, Y., Ostrovski, G., Cain, A., King, H., Summerfield, C., Blunsom, P., Kavukcuoglu, K., Hassabis, D.: Hybrid computing using a neural network with dynamic external memory. *Nature* (2016)
21. Han, X., Wu, Z., Jiang, Y.G., Davis, L.S.: Learning fashion compatibility with bidirectional lstms. In: ACM MM. MM '17 (2017)
 22. Yazici, V.O., Gonzalez-Garcia, A., Ramisa, A., Twardowski, B., van de Weijer, J.: Orderless recurrent models for multi-label classification (2019)
 23. Zaheer, M., Kottur, S., Ravanbakhsh, S., Póczos, B., Salakhutdinov, R.R., Smola, A.J.: Deep sets. In Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R., eds.: *Advances in Neural Information Processing Systems* 30. Curran Associates, Inc. (2017) 3391–3401
 24. Pabbaraju, C., Jain, P.: Learning functions over sets via permutation adversarial networks. CoRR [abs/1907.05638](https://arxiv.org/abs/1907.05638) (2019)
 25. Ramon, J., De Raedt, L.: Multiple instance neural networks. In: *ICML Workshop on Attribute-value and Relational Learning*. (2000)
 26. Wang, X., Yan, Y., Tang, P., Bai, X., Liu, W.: Revisiting multiple instance neural networks. *Pattern Recognition* **74** (2018) 15–24
 27. Liu, Y., Yan, J., Ouyang, W.: Quality aware network for set to set recognition. (2017)
 28. Yan, Y., Wang, X., Guo, X., Fang, J., Liu, W., Huang, J.: Deep multi-instance learning with dynamic pooling. In: *Asian Conference on Machine Learning*. (2018)
 29. Tibo, A., Jaeger, M., Frasconi, P.: Learning and interpreting multi-multi-instance learning networks. [arXiv:1810.11514](https://arxiv.org/abs/1810.11514) (2018)
 30. Tu, M., Huang, J., He, X., Zhou, B.: Multiple instance learning with graph neural networks. In: *ICML Workshop on Learning and Reasoning with Graph-Structured Representations*. (2019)
 31. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. [arXiv:1409.1556](https://arxiv.org/abs/1409.1556) (2014)
 32. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *CVPR*. (2015)
 33. Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed representations of words and phrases and their compositionality. In Burges, C.J.C., Bottou, L., Welling, M., Ghahramani, Z., Weinberger, K.Q., eds.: *Advances in Neural Information Processing Systems* 26. Curran Associates, Inc. (2013) 3111–3119
 34. Devlin, J., Chang, M., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. In: [arXiv:1810.04805](https://arxiv.org/abs/1810.04805). (2018)
 35. Fernando, B., Gavves, E., Oramas M., J., Ghodrati, A., Tuytelaars, T.: Rank pooling for action recognition. In: *TPAMI*. (2016)
 36. Bilen, H., Fernando, B., Gavves, E., Vedaldi, A.: Action recognition with dynamic image networks. *TPAMI* **40** (2018) 2799–2813
 37. Hjelm, R.D., Fedorov, A., Lavoie-Marchildon, S., Grewal, K., Bachman, P., Trischler, A., Bengio, Y.: Learning deep representations by mutual information estimation and maximization. In: *International Conference on Learning Representations*. (2019)
 38. Yang, X., Deng, C., Zheng, F., Yan, J., Liu, W.: Deep spectral clustering using dual autoencoder network. *CVPR* (2019)
 39. Suzgun, M., Gehrmann, S., Belinkov, Y., Shieber, S.M.: LSTM networks can perform dynamic counting. *ACL 2019 Workshop on Deep Learning and Formal Languages* (2019)

40. Liu, Z., Luo, P., Qiu, S., Wang, X., Tang, X.: Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2016)
41. Arandjelović, R., Zisserman, A.: Multiple queries for large scale specific object retrieval. In: BMVC. (2012)
42. Ricci-Vitiani, L., Lombardi, D.G., Pilozzi, E., Biffoni, M., Todaro, M., Peschle, C., De Maria, R.: Identification and expansion of human colon-cancer-initiating cells. *Nature* **445** (2007) 111–115
43. Sirinukunwattana, K., Raza, S.E.A., Tsang, Y., Snead, D.R.J., Cree, I.A., Rajpoot, N.M.: Locality sensitive deep learning for detection and classification of nuclei in routine colon cancer histology images. *IEEE Transactions on Medical Imaging* (2016)
44. Liang, X., Lee, L., Dai, W., Xing, E.P.: Dual motion gan for future-flow embedded video prediction. In: International Conference on Computer Vision. Volume abs/1708.00284. (2017) 1762–1770