Learning to Increment a Contextual Model

Fethiye Irmak Doğan
Dept. of Computer Eng.
METU, Ankara, Turkey
irmak.dogan@metu.edu.tr

Hande Çelikkanat
Dept. of Computer Sci.,
Uni. Helsinki, Finland
celikkan@helsinki.fi

İlker Bozcan, Sinan Kalkan
Dept. of Computer Eng.
METU, Ankara, Turkey
e228998, skalkan@metu.edu.tr

1 Introduction

For a developing agent, learning new skills or representations incrementally from new experiences is of paramount importance. For artificial systems, this problem, named as lifelong learning, continual learning, or incremental learning, has been studied for several decades (see, e.g. [7] for a review), focusing mainly on how to adapt a model when a new class (task) is added to the learning problem owing to a new experience. The established scenario in such studies (e.g., [6, 8]) is to assume that the data is labeled and the agent is aware that the new experience belongs to a different class.

An equally important challenge in lifelong learning is to be able to adapt the model in cases where the class (task) information is not directly observed, i.e., latent. This is especially relevant for an embodied agent that is expected to function in different situations, i.e., contexts, and adapt its behaviors, skills and knowledge to the new contexts. Since context, which can be defined as any prominent configuration of spatial, temporal, social, or task information, can be explicit (directly referred to with labels such as "kitchen", "dinner with guests") or implicit (not directly referred with a label but crucial enough so that the agent should adapt its behavior accordingly), learning the latent context information incrementally is a necessity for a developing agent.

In this paper, we summarize our efforts on incremental learning of latent context information: (i) Celikkanat et al. [1], where we introduced an incremental version of Latent Dirichlet Allocation (LDA) such that the number of contexts (topics) can be incremented automatically with new experience. (ii) Dogan et al. [3], where we modified Restricted Boltzmann Machines to incrementally add new latent variables or layers with new experience. (iii) Dogan et al. [2], where we formulated incrementality of latent contextual information as a learning problem.

In the first two studies [1, 3], our methods for determining when to increment were rule-based, whereas in [2] we showed that incrementing can be learned. Our efforts are different from topic models that try to estimate the optimal number of topics, such as Hierarchical Dirichlet Processes [9], Chinese Restaurant Process and its nested versions [5], and Indian Buffet process [4] in that these methods either assume that all data is available at the estimation phase, or are rule-based.

2 Methods

In this section, we briefly summarize the three methods on incremental learning of latent contextual information. In all methods, objects in the scene are considered as words in a topic model, a scene is taken as a document, and the topics of the document then correspond to the contextual information among the objects in the scene.

2.1 Rule-based Incremental LDA – [1]

In our rule-based incremental LDA model, we start with only one context (i.e. $k = 1$). To determine when to increment k, $\mathbb{C}_{low}$ is specified as a set of objects which have lower confidences than the threshold τ according to their contextual assignments. When there are such objects in $\mathbb{C}_{low}$, i.e., $\mathbb{C}_{low} \neq \emptyset$, the model increments the number of contexts ($k \leftarrow k + 1$) to increase the confidences of the objects in $\mathbb{C}_{low}$.

32nd Conference on Neural Information Processing Systems (NIPS 2018), Montréal, Canada.

2.2 Rule-based Incremental RBM – [3]

Here we summarize our work on (i) incremental Restricted Boltzmann Machines (iRBM) and (ii) deep incremental Boltzmann Machines (diBM).

2.2.1 Incremental Restricted Boltzmann Machines (iRBM)

Our iRBM model depends on the confidence values of each object, which is determined by the weights between the objects (visible neurons) and contexts (hidden neurons):

$$c_v \leftarrow \max_j w_{vj}, \tag{1}$$

where w_{vj} represents the weight between visible unit v and hidden unit indexed by j. In other words, c_v shows the strength of the connection between v and the hidden neurons. If c_v, i.e., the maximum weight to hidden nodes, is low, then v is not strongly represented by any of the current contexts. Moreover, a *baseline* confidence $c_m^{|\mathbf{h}|}$ is also calculated for the entire model:

$$c_m^{|\mathbf{h}|} \leftarrow \frac{1}{Z_0} \exp\left(\min_v c_v\right), \tag{2}$$

with $Z_0 \leftarrow \sum_v \exp(c_v)$ being the partition function. While encountering new scenes, the current confidence ($c_m^{curr} \leftarrow 1/Z_0 \exp\left(\min_v c_v\right)$) can decrease if the objects in the scenes are not strongly represented by any of the current contexts. In order to raise the competence of the model, a new hidden node is added when c_m^{curr} drifts away from $c_m^{|\mathbf{h}|}$:

$$c_m^{curr} < t^{iRBM} \times c_m^{|\mathbf{h}|}, \tag{3}$$

where t^{iRBM}, a scaling factor, adjusts the patience of the model. Furthermore, the weights of the new hidden node (indexed by k) are initialized inversely to the sum of the weights between visible and other hidden nodes, i.e., $w_{ik} \leftarrow 1/\left(\sum_{j=1}^{|\mathbf{h}|-1} w^{ij}\right)$. This weight initialization forces the new topic, i.e., new hidden neuron, to have higher weights to the poorly represented objects and lower weights to the strongly represented ones.

2.2.2 Deep Incremental Boltzmann Machines (diBM)

iRBM increments the number of contexts only for one hidden layer, but in our diBM model, similar contexts in a layer can be represented with a new context in an upper layer of the hierarchy. In order to decide when to increment the number of hidden layers, diBM calculates another baseline confidence, r_f, for the final hidden layer f when the number of hidden neurons in the final layer is exactly two:

$$r_f \leftarrow d(h_i, h_j), \qquad \text{for } h_i, h_j \in \mathbf{h}^f, \tag{4}$$

where $d(h_i, h_j)$ denotes the distance between two hidden nodes, h_i and h_j, depending on their weights:

$$d(h_i, h_j) = \frac{1}{2}\left[D_{KL}(sm(\mathbf{w}^i)||sm(\mathbf{w}^j)) + D_{KL}(sm(\mathbf{w}^j)||sm(\mathbf{w}^i))\right], \tag{5}$$

where $D_{KL}(\cdot||\cdot)$ represents the Kullback-Leibler divergence; $\mathbf{w}^j =< w_{kj} >$ denotes the vector of weights between h_j and the nodes in previous layer; and $sm(\mathbf{w})_i = \exp(w_i)/\sum_j \exp(w_j)$ stands for the vector-defined softmax function. While encountering new scenes and adding new nodes to the layer f, the current confidence of the layer, i.e., $r_f^{curr} \leftarrow \min_{h_i, h_j \in \mathbf{h}^f} d(h_i, h_j)$, drifts away from r_f. In other words, the distance between contexts in layer f become smaller. In order to represent these similar contexts, a new context (hidden) layer with one hidden neuron is added on top of final layer when the following condition satisfies:

$$r_f^{curr} < t^{diBM} \times r_f, \tag{6}$$

where t^{diBM} adjusts the model's patience to increment the number of layers. The weights are initialized randomly between the single node in the newly added layer and nodes in the previous layer.

2.3 Learning-based Incremental LDA – [2]

In our last model, we posed determining when to add a new context as a learning problem (Figure 1), rather than using rules. The main challenge in such an approach is not having a dataset with the correct number of contexts. To address this, we used LDA, which, being a generative model, allows one to sample artificial data with various number of contexts (topics). Since the distribution between the contexts and the objects is important for us, the data being artificial does not pose a problem. The only assumption made here is that contexts in real environments follow a Dirichlet distribution, which, we argue in [2], is a reasonable assumption.

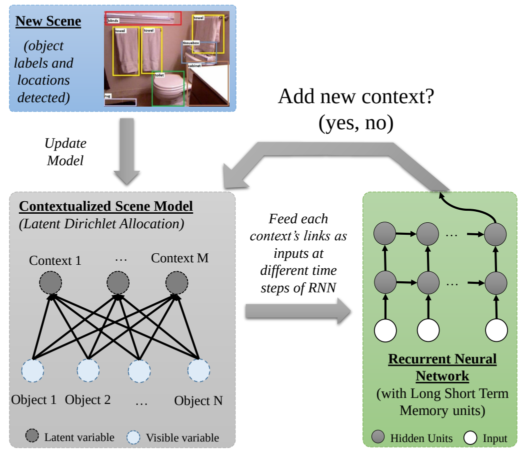

Figure 1: Incrementing context as a learning problem: When a new scene is encountered, LDA Model is provided with the labels of the detected objects. The updated model is fed as input to an RNN, which predicts whether to increment the number of contexts or not.

Assume $D^k = \{S_1^k, ..., S_{L_k}^k\}$ denotes the dataset of scenes generated with k number of contexts. LDA models are trained with k_0 contexts s.t. $k_0 \leq k$. The input-output pair $\{\mathbf{x}, y\}$ used for training the deep network can be constructed as follows:

$\mathbf{x}$: The input vector to the network, describing the current state of the LDA model. Since the number of contexts (hidden topics) is not fixed beforehand, $\mathbf{x}$ should have variable length. Therefore, we take $\mathbf{x}$ to be a sequence of sub-vectors $\mathbf{x}_i = \mathbf{p}_{c_i} = \{p(c_i|o_j)\}_{j=1}^{O}$, i.e., a sequence of conditional probabilities of each context given an object.

y: The expected output of the network; a binary variable (0 or 1) describing whether to increment ($y = 1$) the number of contexts or not ($y = 0$).

2.3.1 Context Incrementing Network (CINet)

Since the input vectors $\mathbf{x}$ have varying lengths, Recurrent Neural Networks (RNNs) are used to cast incrementing as a learning problem, where, given $\mathbf{x}$ describing the state of the LDA model, y (whether to increment the number of contexts) is predicted. Since this is a binary classification problem (increment, no-increment), binary cross-entropy loss is used. As a precaution against over-fitting, L2 regularization on the weights is added.

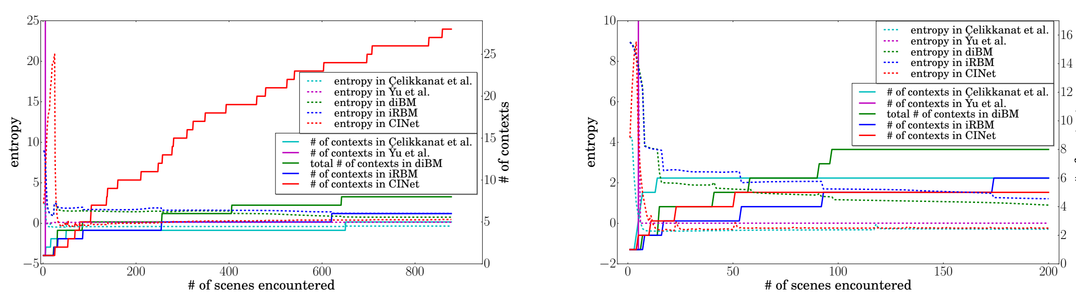

(a) artificial dataset

(b) real dataset

Figure 2: How the different methods increment the number of contexts, and how the entropy of the system changes on the artificial dataset (a) and the real dataset (b). In (a), the data has 5 contexts whereas in (b) it has 25. Note that CINet was trained only on the artificial data.

3 Results

We now present some of the results we have obtained and reported in [2]. In Figure 2, we see our three methods as well as another study [10] applied on the artificial data generated by LDA as well as on a scene classification dataset (SUN RGB-D) where we carefully selected a subset (with clear contexts). We see that the learning-based approach outperforms other methods in terms of the number contexts it converges to, while maintaining a low entropy. The result on the real dataset is striking since the model was trained on the artificial data. This suggests that the word-topic distribution assumed by the data generating LDA matches the context-object distribution in real settings. We want to elaborate more on the generalization capability of the learning-based method. In Figure 3, we see the probabilities of incrementing suggested by CINet in different settings (with 7 or 15 contexts as the ground truth). See [2] for more results. We did not experiment with topic modeling datasets since a document can have more than one topic, which makes it difficult to assess the converged number of contexts of our models.

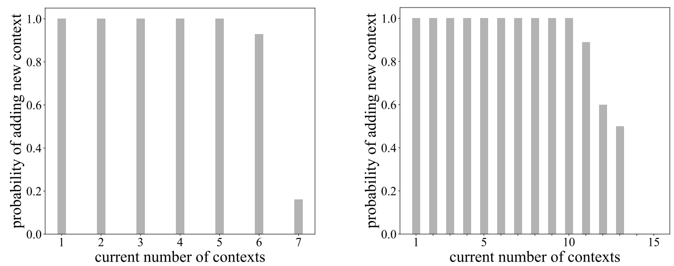

(a) (b)

Figure 3: Probability of incrementing contexts for various states of an LDA model on the artificial data. Ground truth is respectively (a) 7, and (b) 15. Note that the network was trained for LDA models up to 10 contexts.

4 Summary

In this paper, we summarized our efforts on incremental construction of latent variables in context (topic) models. With our models, an agent can incrementally learn a representation of critical contextual information. We demonstrated that a learning-based formulation outperforms rule-based models, and generalizes well across many settings and to real data.

Acknowledgments

This work was supported by the Scientific and Technological Research Council of Turkey (TÜBİTAK) through project called "Context in Robots" (project no 215E133).

References

[1] H. Celikkanat, G. Orhan, N. Pugeault, F. Guerin, E. Şahin, and S. Kalkan. Learning context on a humanoid robot using incremental latent dirichlet allocation. *IEEE TCDS*, 8(1):42–59, 2016.

[2] F. I. Dogan, I. Bozcan, M. Celik, and S. Kalkan. Cinet: A learning based approach to incremental context modeling in robots. *IEEE/RSJ IROS*, 2018.

[3] F. I. Dogan, H. Celikkanat, and S. Kalkan. A deep incremental boltzmann machine for modeling context in robots. *IEEE ICRA*, 2018.

[4] T. L. Griffiths and Z. Ghahramani. The indian buffet process: An introduction and review. *JMLR*, 12:1185–1224, 2011.

[5] T. L. Griffiths, M. I. Jordan, J. B. Tenenbaum, and D. M. Blei. Hierarchical topic models and the nested chinese restaurant process. In *NIPS*, 2004.

[6] J. Kirkpatrick, R. Pascanu, N. Rabinowitz, et al. Overcoming catastrophic forgetting in neural networks. *PNAS*, 2017.

[7] G. I. Parisi, R. Kemker, J. L. Part, C. Kanan, and S. Wermter. Continual lifelong learning with neural networks: A review. *arXiv preprint arXiv:1802.07569*, 2018.

[8] A. A. Rusu, N. C. Rabinowitz, G. Desjardins, et al. Progressive neural networks. *arXiv preprint arXiv:1606.04671*, 2016.

[9] Y. W. Teh, M. I. Jordan, M. J. Beal, and D. M. Blei. Sharing clusters among related groups: Hierarchical dirichlet processes. In *NIPS*, 2005.

[10] J. Yu, J. Gwak, S. Lee, and M. Jeon. An incremental learning approach for restricted boltzmann machines. In *ICCAIS*, 2015.