Stanford University / CS231N / Spring 2017

CHILDNet: Curiosity-driven Human-In-the-Loop Deep Network

Byungwoo Kang[1], Hyun Sik Kim[2], Donsuk Lee[3]

INTRODUCTION

Humans can learn actively and incrementally

Active Learning

- Human annotation is expensive
- Choose examples to request labels for

Incremental Learning[5]

- New visual concepts emerge in the real world
- Learn new concepts while preserving existing knowledge

➢ **Active + Incremental Learning System**

PROBLEM STATEMENT

Goal

- Class-incremental online active learning from continuous unlabeled image stream

Task Definition

- Stream of unlabeled image data
- Labels are only available upon request
- Number of classes can increase over time
- Learn to recognize new classes online using as few examples as possible

Evaluation

- Label request rate
- Prediction Accuracy

DATASET

Omniglot[1]

- 1,623 classes of characters
- 20 hand drawn examples per class
- Total 32,460 examples
- 800 classes for training, 400 classes for validation, 423 classes for testing

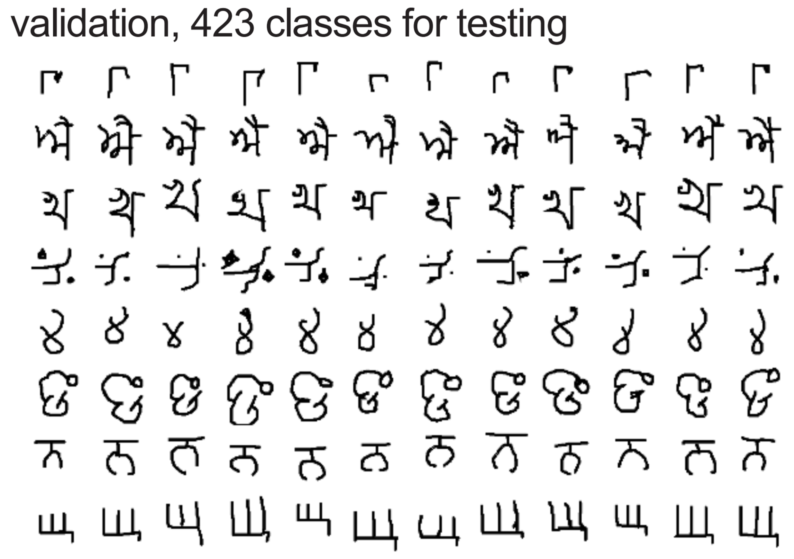

METHODS

Siamese Net for One-shot Classification[7]

- Trained for same/different pair verification tasks
- Maintains running mean of features for each class
- Given input image x, predict class C^* such that

$$C^* = argmin_c \boldsymbol{p}^{(c)}(\boldsymbol{x})$$

where $\boldsymbol{p}^{(c)}(\boldsymbol{x})$ is similarity score between x and mean of class c

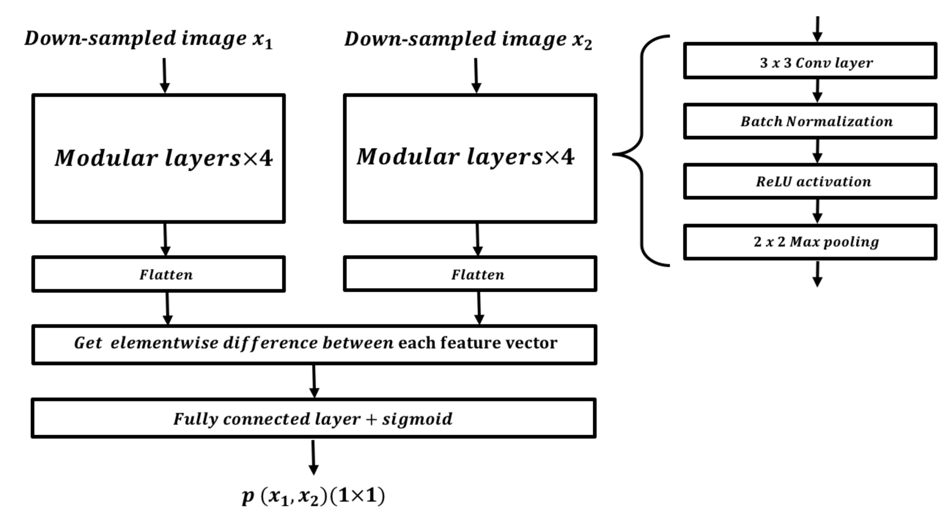

Reinforcement Learning Model

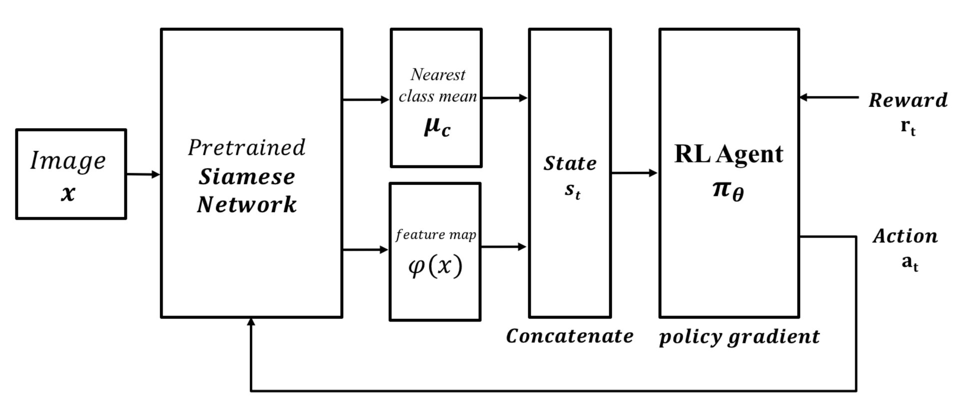

$$Reward \begin{cases} R_{cor}\text{, if predicting and prediction correct} \\ R_{inc}\text{, if predicting and prediction incorrect} \\ R_{req}\text{, if requesting a label} \end{cases}$$

- Siamese network for feature extraction and incremental learning
- RL agent decides whether to request a label or make a prediction
- Policy gradient with parameter update rule:

$$\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} + \boldsymbol{\alpha}\, \nabla_{\boldsymbol{\theta}} \boldsymbol{log}\,(\boldsymbol{\pi_\theta}\,(\boldsymbol{s_t}, \boldsymbol{a_t}))\boldsymbol{v_t}$$

where, v_t is an episode discounted reward

EXPERIMENTS

Train

- Sample 10 classes per episode
- 30 images from the 10 classes per episode
- 20,000 episodes for training

Test

- N classes per episode
- All images from N classes per episode

RESULTS

Label Requests

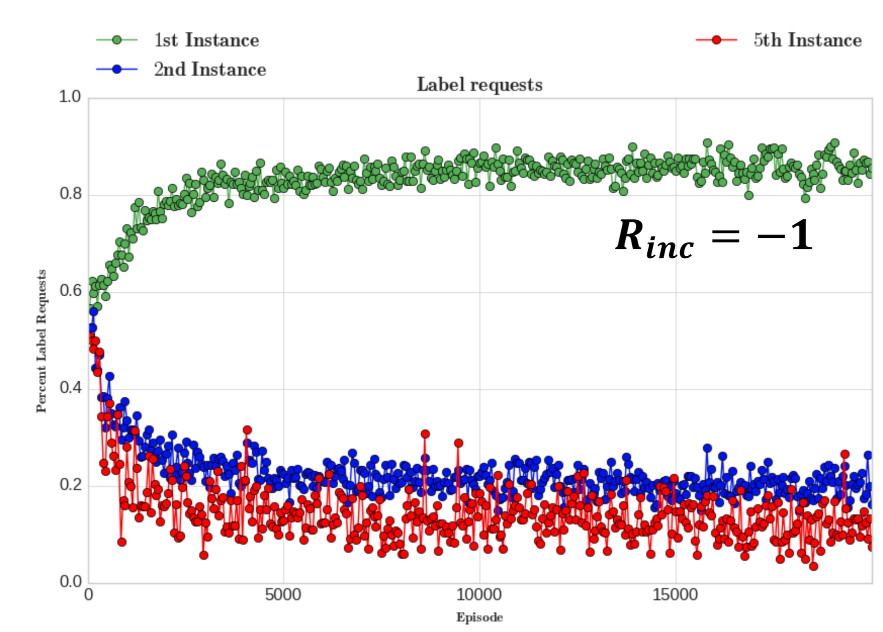

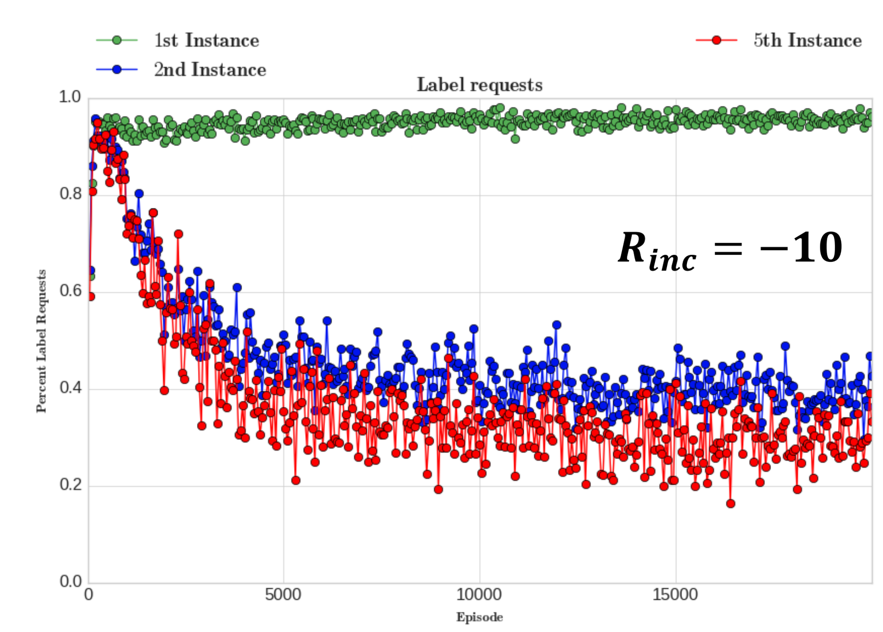

Prediction Accuracy

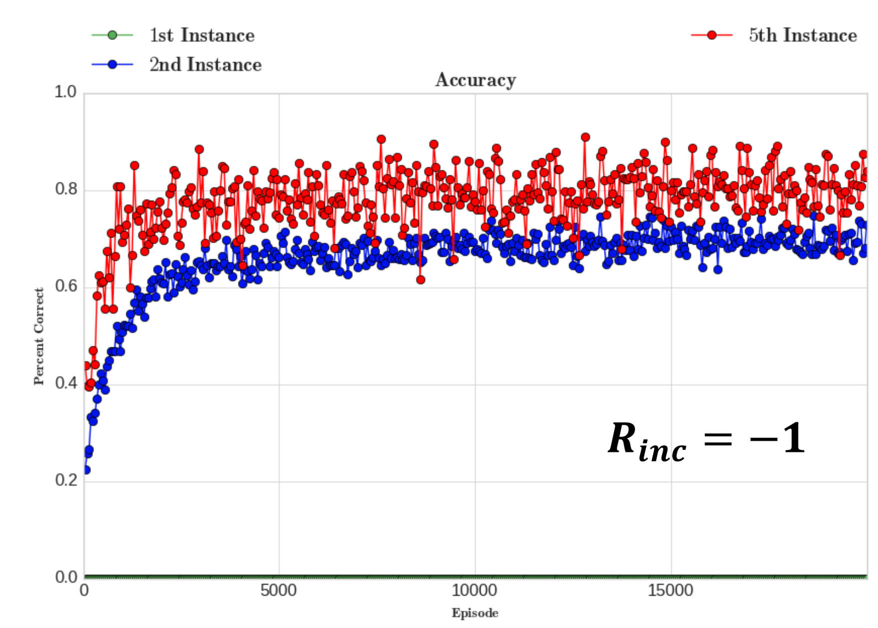

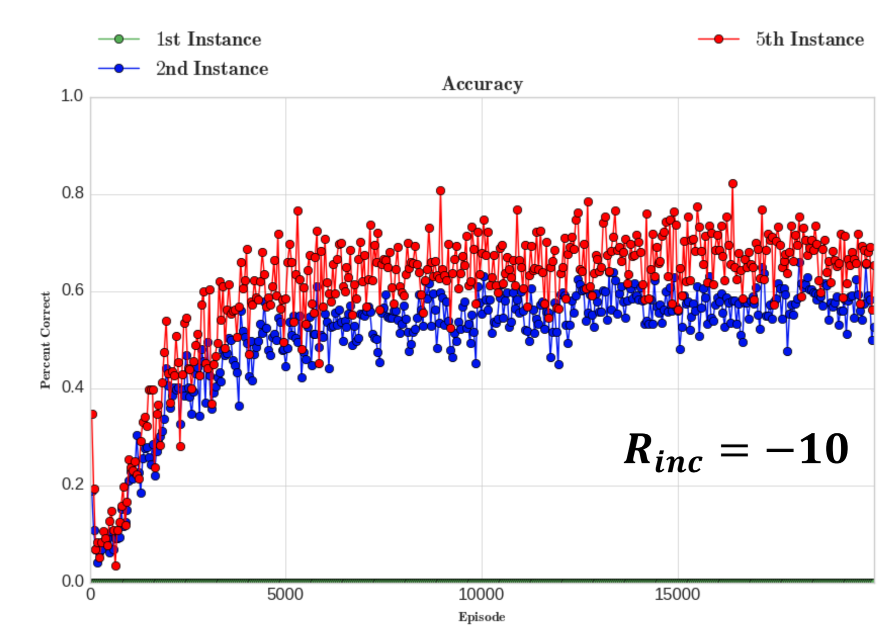

- The model makes fewer requests and achieves higher accuracy on later instances of a class
- It makes decisions based on uncertainty about its own knowledge

Trading Accuracy for Requests

	Accuracy(%)	Requests(%)
Supervised	93.4	100
RL Prediction($R_{inc} = -1$)	87.7	17.0
RL Prediction($R_{inc} = -5$)	92.1	24.1
RL Prediction($R_{inc} = -10$)	93.2	26.3

- % of correct predictions and % of label requests
- Increasing the penalty for an incorrect prediction improves accuracy at the cost of more label requests

Varying number of classes

	Accuracy(%)	Requests(%)
Number of classes = 3	96.4	16.7
Number of classes = 10	87.7	17.0
Number of classes = 20	77.1	17.2
Number of classes = 40	64.9	16.7

- The model is trained with 10 classes per episode
- The model is applicable to variable number of classes

CONCLUSION

- Class-incremental[5] online active learning is possible
- The choice of rewards can trade off accuracy for requests

FUTURE WORK

- Experiment the model with larger datasets, e.g. ImageNet
- Develop automatic data annotator with humans in the loop

BIBLIOGRAPHY

1 S. Ager. Omniglot - writing systems and languages of the world. Inwww.omniglot.com, 2015.
2 C. Finn, P. Abbeel, and S. Levine. Model-agnostic meta-learning for fast adaptation of deep networks, 2017. arXiv preprint arXiv:1703.03400.
3 K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition, 2016. In Computer Vision and Pattern Recognition (CVPR).
4 E. Lughofer. Single-pass active learning with conflict and ignorance, 2012. In Evolving Systems.
5 S. Rebuffi, A. Kolesnikov, G. Sperl, and C. H. Lampert. iCaRL: Incremental classifier and representation learning, 2017. In Computer Vision and Pattern Recognition (CVPR).
6 M. Woodward and C. Finn. Active one-shot learning, 2016. NIPS.
7 G. Koch, R. Zemel, and R. Salakhutdinov. Siamese Neural Networks for One-shot Image Recognition, in 2015. In 32nd International Conference on Machine Learning (ICML).

Stanford University

1. bkang@stanford.edu
2. hsik@stanford.edu
3. donlee90@stanford.edu