RIT | Global Cybersecurity Institute

Cross-Organizational Continual Learning of Cyber Threat Models

Chanel Cheng, Shanchieh Jay Yang

Introduction

- **Intrusion detection systems are developed to help detect cyber threats across networks.**
- **Yet, cyber threats evolve over time and follow different patterns across various organizations.**
  - **Continuous detection of changing data is difficult by traditional means[1]**

Consider:

An incoming stream of network traffic from two different organizations.

- **Similar attack types have different patterns across organizations**
- **New attack types are also present in each organization**

Stream encounters both gradual and sudden changes in attack patterns

Related Works

PNNs / EWC / SI / iCaRL / GEM

- **PNNs[2]** - constructs new networks as novel tasks occur, resulting in **linearly increasing memory requirement.**
- **EWC & SI[3,4]** - extends loss function with a term that consolidates selective network weights, but requires **explicit task boundaries.**
- **iCaRL[5]** - combined use of replay and distillation but still requires **explicit task boundaries.**
- **GEM[6]** - builds optimization constraints using old data but **less effective across shifting domains.**

Methods	PNN [35]	PackNet [28]	HAT [37]	ER [31, 33]	MER [33]	GSS [1]	GEM [27]	A-GEM [9]	HAL [8]	iCaRL [32]	FDR [4]	LwF [24]	SI [42]	oEWC [20]	**DER (ours)**	**DER++ (ours)**
Constant memory	–	–	–	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
No task boundaries	–	–	–	✓	✓	✓	–	✓	–	–	–	–	–	–	✓	✓
No test time oracle	–	–	–	✓	✓	✓	✓	✓	✓	✓	✓	–	✓	✓	✓	✓

Table.1. Side-by-side comparison of continual learning strategies[6]

Experience Replay (ER)

- **ER[7]** eliminates need for task boundaries, test time oracle, and enforces constant memory footprint.
- **Potentially more suited for real-world scenarios with gradual and sudden shifts in data.**

Methodology

Table.2. Aggregate port mapping and one-hot encoding maps the most commonly used port numbers to their corresponding port services and one-hot encodes them as features for the model to learn from.

Port Service	*Port #*
DNS	53
http	80, 8080
https	443, 8443
wbt	3389
smb	445, 139, 137
ftp	20, 21
ssh	22
llmnr	5535
other	(unassigned port #'s)

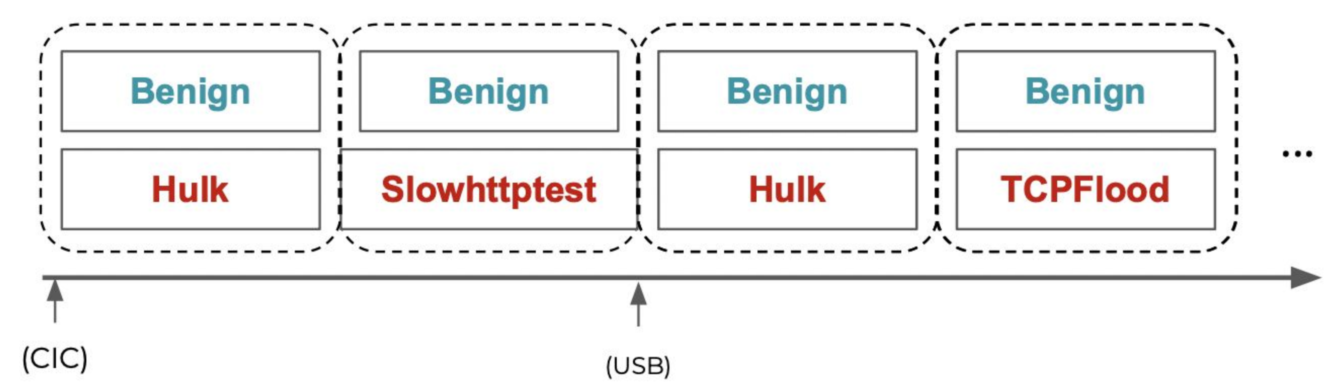

Figure.1. Example data stream for task-agnostic continual learning on network traffic flow from CIC-IDS-2018[8] and USB-IDS-2021[9]

- **Two datasets were converted into a single data stream for continual learning without task boundaries.**
  - **Regular benign traffic and malicious traffic are present together in stream (with benign as the majority of traffic)**
  - **Order and source of data does not matter for our continual model in learning the attack types**

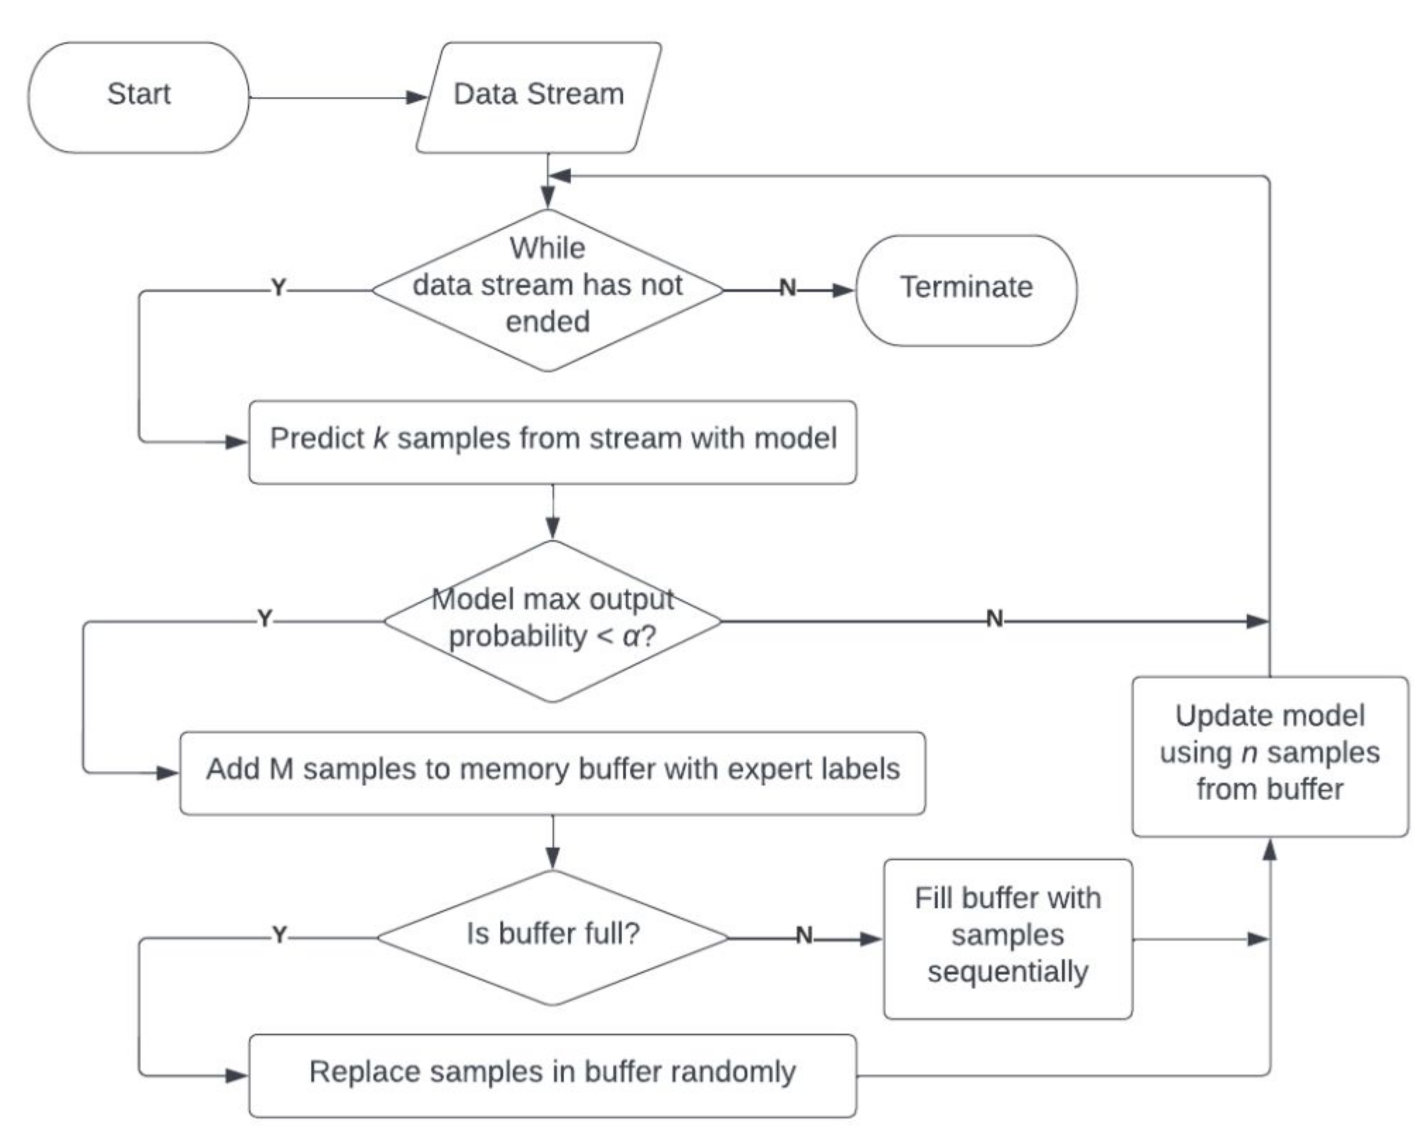

Figure.2. Flow diagram for the continual learning strategy.

- **Replay buffer of fixed size with older samples replaced as new samples are selected to the buffer.**
- **Only expert-labeled samples saved to the buffer train and update the model.**

Experiment & Results

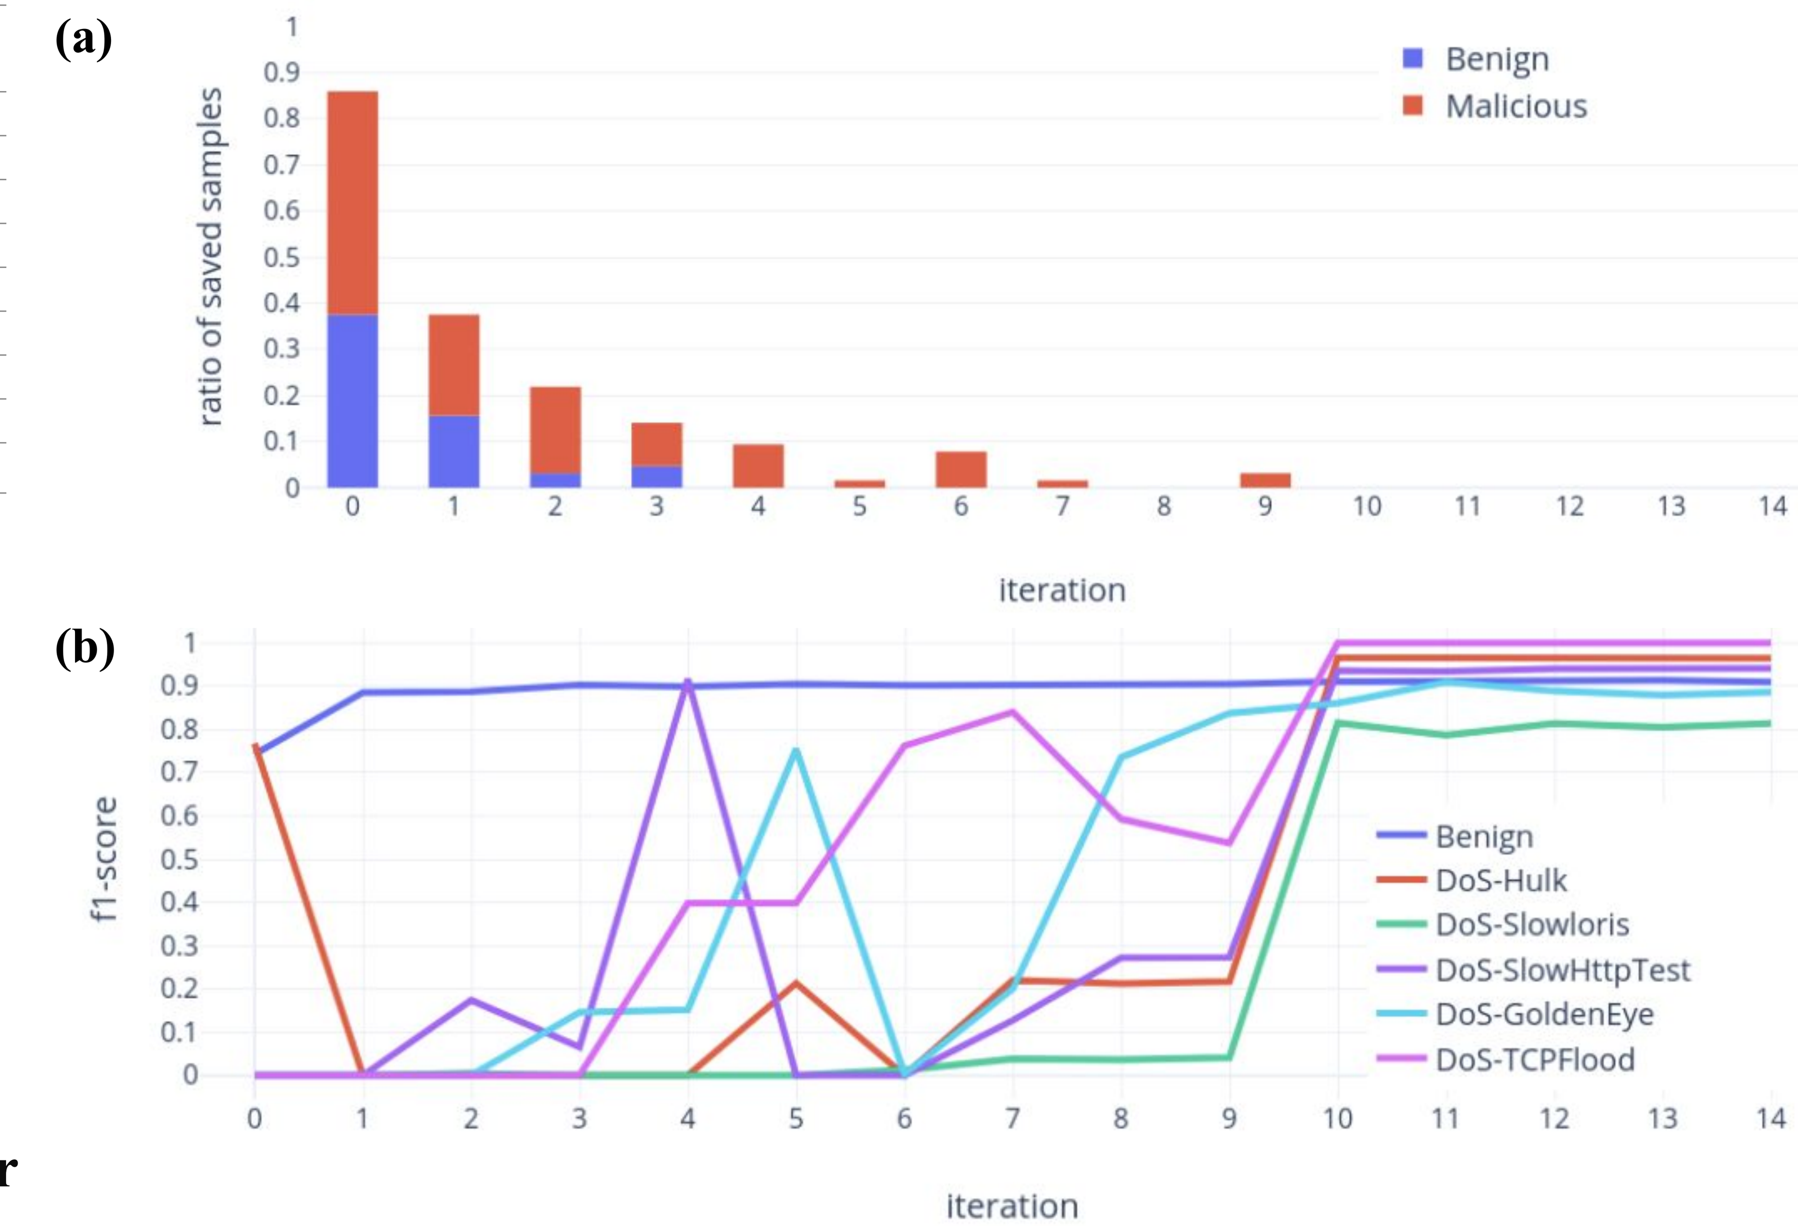

Figure.3. (a) Average ratio of samples saved to buffer and (b) f1-score for network traffic classification as new DoS classes were introduced over 14 iterations.

Observations & Future Works

- By the 10th iteration, no more samples are required to be labeled, while most of the F1-scores reach above 0.9 except DoS-GoldenEye (~0.88) and DoS-Slowloris (~0.8), both of which have a much smaller sample size.
- This learning strategy largely reduces the number of labeled data needed and can quickly reach good prediction performance.

Future Works:

1) Expand the experiment to include more attack types, 2) further investigate class imbalance issue, and 3) optimize the sampling strategy for replay buffer.

References

[1] D. Silver, Q. Yang, and L. Li. 2013. Lifelong machine learning systems: Beyond learning algorithms. In 2013 AAAI spring symposium series.

[2] A. Rusu, N. Rabinowitz, G. Desjardins, H. Soyer, J. Kirkpatrick, K. Kavukcuoglu, R. Pascanu, and R. Hadsell. 2016. Progressive neural networks. arXiv preprint arXiv:1606.04671 (2016).

[3] J. Kirkpatrick, R. Pascanu, N. Rabinowitz, J. Veness, G. Desjardins, A. Rusu, K. Milan, J. Quan, T. Ramalho, A. Grabska-Barwinska, et al. 2017. Overcoming catastrophic forgetting in neural networks. Proceedings of the national academy of sciences 114, 13 (2017), 3521–3526.

[4] F. Zenke, B. Poole, and S. Ganguli. 2017. Continual learning through synaptic intelligence. In International Conference on Machine Learning. PMLR, 3987–3995

[5] S.-A. Rebuffi, A. Kolesnikov, G. Sperl, and C. Lampert. 2017. icarl: Incremental classifier and representation learning. In Proceedings of the IEEE conference on Computer Vision and Pattern Recognition. 2001–2010.

[6] D. Lopez-Paz and M. Ranzato. 2017. Gradient episodic memory for continual learning. Advances in neural information processing systems 30 (2017).

[7] P. Buzzega, M. Boschini, A. Porrello, D. Abati, and S. Calderara. 2020. Dark experience for general continual learning: a strong, simple baseline. Advances in neural information processing systems 33 (2020), 15920–15930.

[8] "CIC-IDS-2018 on AWS." https://www.unb.ca/cic/datasets/ids-2018.html (accessed Jul. 18, 2022).

[9] "USB-IDS-1." http://idsdata.ding.unisannio.it/datasets.html (accessed Jul. 18, 2022).