

Detecting Concept Drift in Fully Distributed Environments

István Hegedűs, Nyers Lehel and Róbert Ormándi

University of Szeged, Szeged, Hungary and
Subotica Tech, Subotica, Serbia

Motivation

- We are in the middle of big data era.
- Our devices and applications accumulate more and more data.
- Paradoxically while getting access to more and more data, we can find less information by applying the original retrieval techniques.
- This phenomena increases the importance of emerging **fully distributed, large-scale data mining** algorithms that can work in **unreliable networks**, take the **privacy** of users into account and **be adaptive** to changes in the data.


Supervised Learning

- We have a manually labeled training database

$$S = \{(x_1, y_1), \dots, (x_l, y_l)\} \in \mathbb{R}^d \times \{-1, +1\}.$$

- Here x_i called **feature vector** that describes an object of real world (email) as a real-valued vector
- y_i is the class label which assigns the feature vector to a well-defined class (e.g. spam).
- The goal—through the learning phase—of the classification problem is to find a model $f : \mathbb{R}^d \rightarrow \{-1, +1\}$ which can classify any sample coming from the same probability distribution (D).

Illustration of classification with a linear model



Fully distributed data

- Horizontal data distribution
- Every node has very few samples, we assume they have **only one**
- We do not allow for moving data, only local processing (**privacy preservation**)
- **Drifting**: data distribution can change

GoLF

- Gossip Learning Framework (GoLF) - is a learning framework designed for performing fully distributed learning in large scale networks.
- The basic idea behind GoLF is that large number of online models take random walks in the network while improving themselves using the training samples contained at the nodes and getting combined by applying ensemble learning techniques.

Concept Drift

- The distribution that generates the training databases (D) may change over the time. This change—especially when it is significant and sudden—makes the already trained models inaccurate. This phenomena is called concept drift.
- we have to identify the time moment t when (sudden) drift occurred while we also handle it. This problem is known as the **drift detection problem**.

Algorithm 1 CDDGoLF

```
1:  $c \leftarrow 0$ 
2:  $\text{currentModel} \leftarrow \text{initModel}()$ 
3:  $\text{receivedModels.add}(\text{currentModel})$ 
4: loop
5:   if  $\text{receivedModels} = \emptyset$  then
6:      $c \leftarrow c + 1$ 
7:   if  $c = 10$  then
8:      $\text{receivedModels.add}(\text{currentModel})$ 
9:   for all  $m \in \text{receivedModels}$  do
10:     $p \leftarrow \text{selectPeer}()$ 
11:     $\text{send } m \text{ to } p$ 
12:     $\text{receivedModels.remove}(m)$ 
13:     $c \leftarrow 0$ 
14:   $\text{wait}(\Delta)$ 

15: procedure ONRECEIVEMODEL( $m$ )
16:   $m \leftarrow \text{driftHandler}(m)$  ▷ Drift
17:   $\text{currentModel} \leftarrow \text{updateModel}(m)$ 
18:   $\text{receivedModels.add}(\text{currentModel})$ 
```

Algorithm 2 Procedures

```
1: procedure INITMODEL
2:    $m.w \leftarrow (0, 0, \dots)^T$ 
3:    $m.age \leftarrow 0$ 
4:    $m.history \leftarrow \emptyset$ 
5:   return  $m$ 

6: procedure DRIFTHANDLER( $m$ )
7:    $\hat{y} \leftarrow m.\text{predict}(x)$ 
8:    $m.history.add(\hat{y} = y ? 0 : 1)$  ▷ history update
9:   if  $\text{driftOccurred}(m)$  then
10:     $m \leftarrow \text{initModel}()$ 
11:   return  $m$ 

12: procedure DRIFTOCCURRED( $m$ )
13:    $\text{errorRates} \leftarrow \text{smooth}(m.history)$ 
14:    $\text{slope} \leftarrow \text{linearReg}(\text{errorRates})$ 
15:   if  $\text{rand}([0; 1]) < \sigma_{c,d}(\text{slope})$  then
16:     return true
17:   return false

18: procedure UPDATEMODEL( $m$ )
19:    $m.age \leftarrow m.age + 1$ 
20:    $\hat{y} \leftarrow m.\text{predict}(x)$  ▷  $(x, y)$  is stored locally
21:    $m.w \leftarrow (1 - \frac{1}{m.age})m.w + \frac{1}{m.age \cdot \lambda}(y - \hat{y})x$ 
22:   return  $m$ 

23: procedure PREDICT( $x$ )
24:    $p_0 \leftarrow 1 / (1 + \exp(\text{currentModel}.w^T x))$ 
25:    $p_1 \leftarrow 1 - p_0$ 
26:   return  $p_0 > p_1 ? 0 : 1$ 
```

The solution – detecting drift

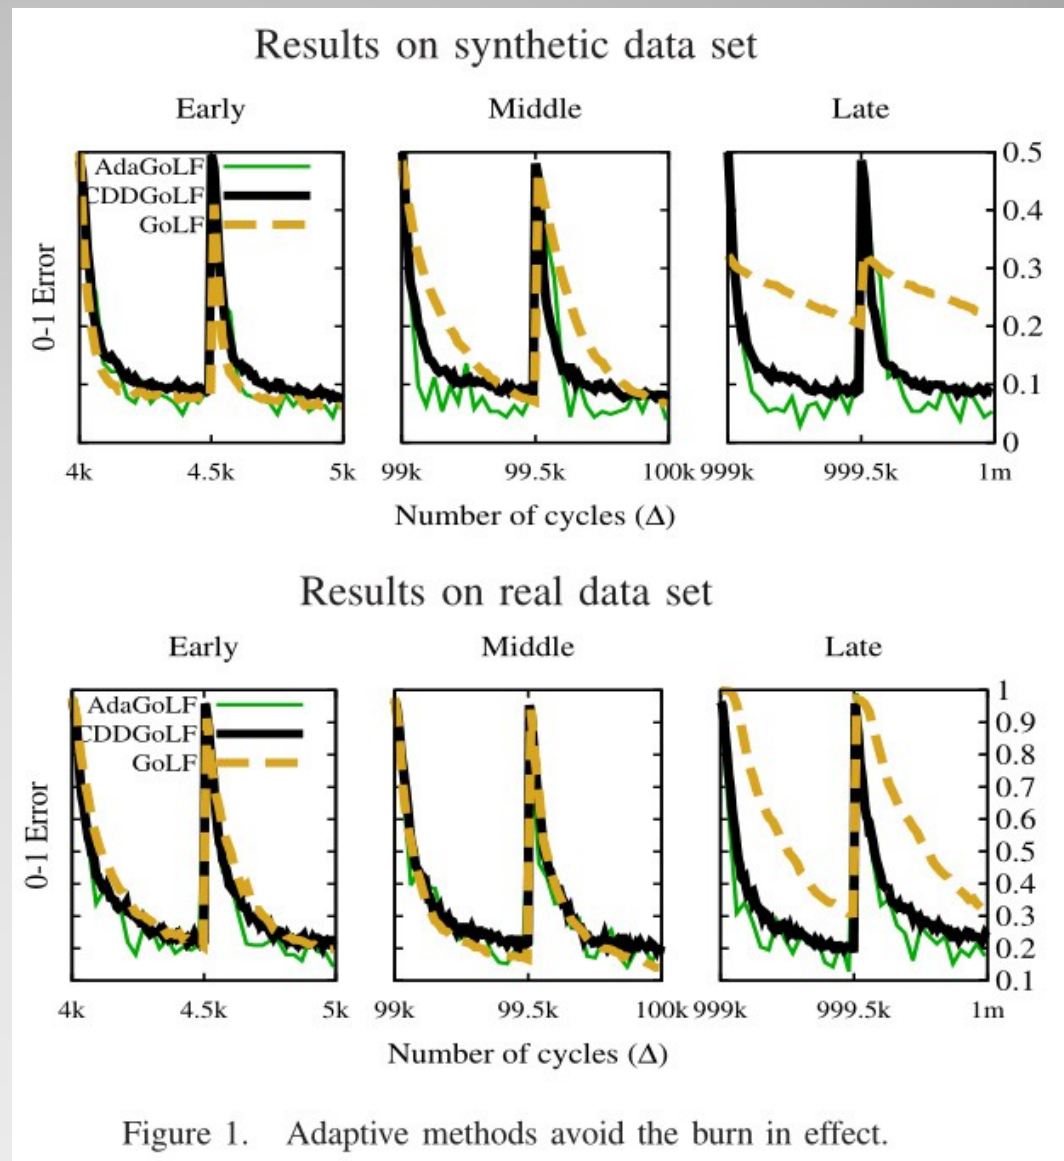
```
6: procedure DRIFTHANDLER(m)
7:    $\hat{y} \leftarrow m.predict(x)$ 
8:   m.history.add( $\hat{y} = y ? 0 : 1$ )    ▷ history update
9:   if driftOccurred(m) then
10:    m ← initModel()
11:   return m

12: procedure DRIFTOCCURRED(m)
13:   errorRates ← smooth(m.history)
14:   slope ← linearReg(errorRates)
15:   if rand([0; 1]) <  $\sigma_{c,d}(slope)$  then
16:    return true
17:   return false
```

- if the slope is negative or 0, the learning is probably in the convergence or in the already converged phase, respectively. Otherwise, if the slope is
- positive, drift probably occurred.

Experimental results: Drift Handling

- For both databases we modeled the drift by changing the labeling over the time.
- Algorithms without concept drift handling capabilities show **burn in** effect, i.e. after a certain time they cannot adapt to the changing concepts.



Experimental results: Drift Detection




Figure 2. The drift detection and the classification performance of the proposed method on synthetic and real datasets.

CONCLUSION

- In this work we proposed a novel algorithm for detecting concept drifts in a fully distributed network on the top of our previously proposed learning framework (GoLF).
- The approach introduces a cache, called history, for collecting data about the performance of the models. Later in the protocol a mechanism decides whether drift occurred.
- Through empirical evaluations we investigated the performance (both for drift handling and detection capabilities) and scalability of the method. Based on these, we pointed out that our current proposal is a suitable choice when we need the additional feature of drift detection, since the drift handling performance of our approach is similar to the state-of-the-art approaches.

Thank you for your time!