Teacher Student Curriculum Learning applied to OCR

Rodrigo Laguna, *supervised by Guillermo Moncecchi*

Universidad de la República - PEDECIBA

Motivation

Curriculum Learning (CL) mimics the natural learning process in humans by presenting examples in a specific order, it was proposed by Elman [3] while the term was coined by Bengio [1]. However, some key aspects, such as measuring task complexity and designing the curriculum, still require standardization. In this study, curriculum design is addressed through Reinforcement Learning (RL). CL is applied to the challenging task of Optical Character Recognition (OCR).

The LUISA project's OCR task was selected because previous attempts to develop custom models for this particular dataset were not able to outperform existing open-source alternatives. By applying CL methods, which theoretically require less training or data, the goal is to improve the performance of OCR on this challenging dataset.

Problem Description

The dataset consists on small images containing text extracted from documents produced between 1958-1985 during last civic-military dictatorship ruled in Uruguay, that were found in the Ministry of Defence in 2006-2007.

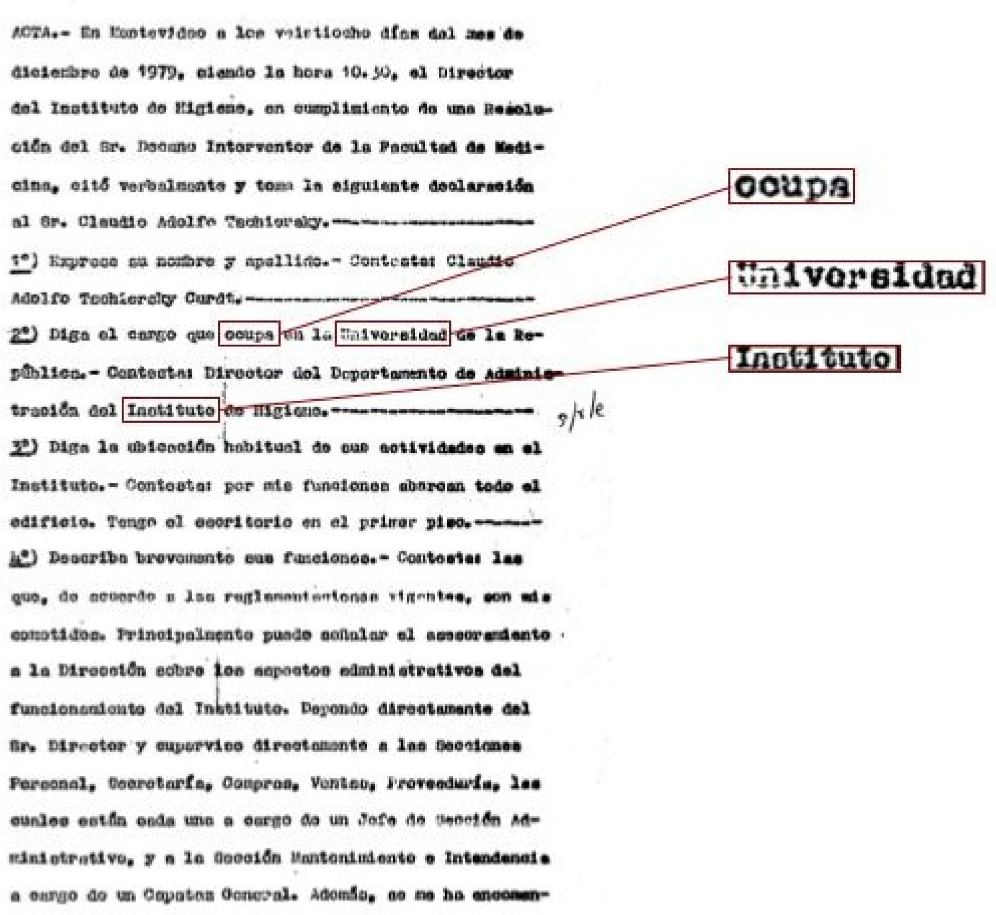

Those examples that could not be well-transcribed with open source OCR tools, were given to volunteers to be labeled as part of the LUISA project, to later train a custom OCR.

Labelled images had several damage caused by the time and bad conservation conditions, and some distortion due to the digitalization process, which may explain why available OCR tools do not perform well on them.

Methodology

The goal is to understand the differences between using or not the TSCL [4] mechanism to train the OCR model, with respect to the usual training process. The dataset, model and hyper-parameters are defined by Chavat [2] and therefore, used as a strong baseline.

Teacher Student Curriculum Learning

In traditional SGD, learning examples are randomly shuffled:

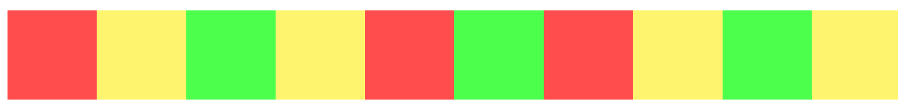

However, humans and animals learn first from easy examples while harder are learned later:

In this work, shorter examples are assumed to be easier than longer ones, in terms of $\#Chars$.

Easy
Medium
Hard

The examples are separated into 20 sub tasks, according to the number of chars on each. Then, there is a Teacher that predicts a distribution function over the sub tasks, which is used to sample the examples used to train the Student during the next step.

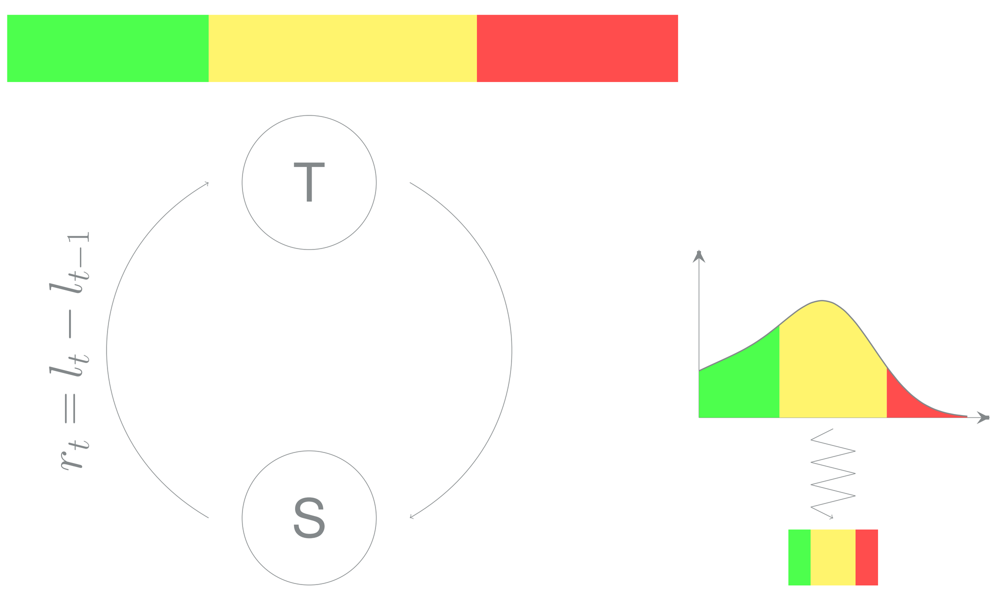

The Student is an OCR model that is fit on each step, only once on each sampled example during the step.

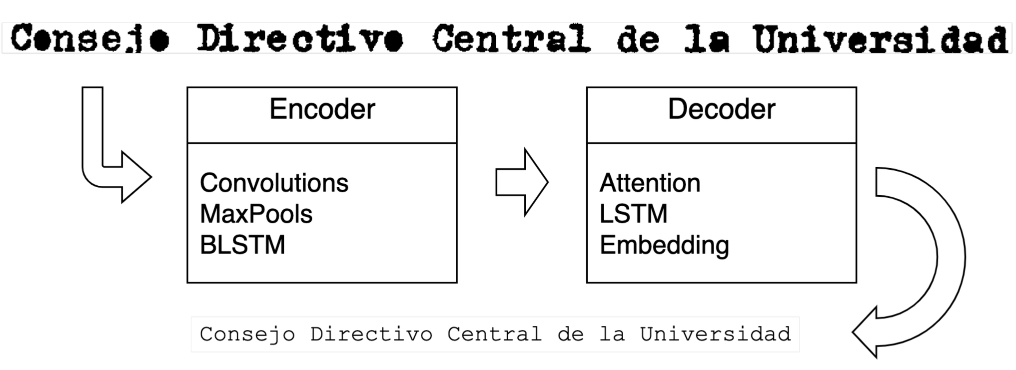

Finally, the loss is computed for each sub task and given as feedback to the Teacher, that is trained using RL, considering as a reward the slope in the loss function between two steps.

The key idea here is to **train more often in those tasks that are being learned (or forgotten) the most**.

Training

Both, the Teacher and Student, are trained at the same time, within a single session.
The OCR model is fit using ADAM algorithm on each step. There is also a teaching forcing mechanism, which consits on randomly feeding the proper output instead of the predicted one to train the decoder.

```
CL = c_learner(train_full, val_full, n_sub_tasks=20)

for e in range(TOTAL_EPOCHS):
    while remaining_steps_on_epoch > 0:
        # sample instances to train
        train_i, val_i = CL.choose_task()

        # fit breaks if  remaining_steps_on_epoch <= 0
        fit_model(train_i)

        loss = eval_model(val_i)

        # Compute reward and update Q function
        CL.learn(loss)

    # epoch ended
    compute_metrics(val_full)
```

Partial Results

Current results seem comparable with the baseline. They were only measured in the validation set, which is also used for early stopping as defined by Chavat [2], and is also used to guide the teacher's training, hence they may be overestimated. That said, they are still comparable among them. A new iteration of experiments and further parameter tuning is still pending, therefore, the test set was never used in this work up to the moment.

Algorithm	Epochs	CER	LCS	Loss
Eps-Greedy	32/32	27.9	74.0	2.946
Softmax	26/32	28.6	73.3	2.864
Random	24/32	29.8	72.1	2.703
Baseline	**16/32**	**27.2**	**74.5**	**2.209**

Conclusion and Future Work

This study of TSCL in OCR task shows promising results, despite not surpassing the baseline yet. The use of TSCL has a significant impact on the training process, and further analysis is needed to understand it.

To contribute with labeling, contact or download a copy of this poster, scan this QR:

References

[1] Yoshua Bengio et al. "Curriculum learning". In: *Proceedings of the 26th annual international conference on machine learning.* 2009, pp. 41–48.
[2] Felipe Chavat Pérez. "Modelos Seq2Seq para la transcripción de documentos del Archivo Berrutti". In: (2022).
[3] Jeffrey L Elman. "Learning and development in neural networks: The importance of starting small". In: *Cognition* 48.1 (1993), pp. 71–99.
[4] Tambet Matiisen et al. "Teacher-student curriculum learning". In: *IEEE transactions on neural networks and learning systems* (2019).

Contact: rodrigo.laguna@fing.edu.uy