

WARC-DL: SCALABLE WEB ARCHIVE PROCESSING FOR DEEP LEARNING

Niklas Deckers Martin Potthast
Leipzig University




Figure 1: The six steps of the WARC-DL web archive processing pipeline: Filters and the Keras training are customizable.

Introduction. Web archives have grown to petabytes. In addition to providing invaluable background knowledge on many social and cultural developments over the last 30 years, they also provide vast amounts of training data for machine learning. To benefit from recent developments in Deep Learning, the use of web archives requires a scalable solution for their processing that supports inference with and training of neural networks. To date, there is no publicly available library for processing web archives in this way, and some existing applications use workarounds [1]. This paper presents WARC-DL,¹ a deep learning-enabled pipeline for web archive processing that scales to petabytes.

Technical Setting. In many “traditional” data centers, mass storage, CPU-bound processing, and GPU-bound processing are separate clusters. Since the storage capacity of the latter two does not usually match that of the former, large web archives are traditionally processed in batches. However, batch processing of web archives on GPU clusters is wasteful: since web archive data is raw data, it must be preprocessed before being passed to a GPU. This preprocessing is usually CPU-bound and highly parallelizable, so using a CPU cluster for this is desirable before streaming the preprocessed data to the GPU cluster. In addition, the data relevant to a particular task is usually sparse across the archive (e.g., only certain images are needed for the training of image representations, and only certain plaintexts are needed for argument mining), resulting in a variable data flow after preprocessing. To optimize GPU usage, a constant flow of data is required. WARC-DL solves both problems: After loading web archive data from the memory cluster into the CPU cluster for preprocessing, it streams the useful data into the GPU cluster for interleaved processing.

Web Archive Processing Pipeline. Figure 1 provides an overview of the six-step pipeline implemented by WARC-DL: (1) WARC files² are distributed to the CPU workers using PySpark. (2) FastWARC [2]³ is used to decompress and iterate the records. The records can optionally be filtered and CPU-bound preprocessing like feature extraction or tokenization can be performed. (3) Pickled record streams are sent to the GPU Cluster via TCP. (4) The streams are converted into TensorFlow datasets and interleaved. (5) One option is to use a pre-trained Keras model to batch process the samples. The results are filtered and stored, including the original data extracted from the WARC datasets. (6) Alternatively, the samples can be used to train a Keras model after an optional filtering step, e.g., for duplicate removal. The filtering in (2) and (6) implements a basic MapReduce.

The preprocessing steps and the pre-trained models used are fully customizable (an extension to frameworks other than TensorFlow is planned). The pipeline enables extraction of data from multiple modalities, including text for language models and images for computer vision models. It also supports the simultaneous extraction of linked web archive records, such as the text of a web page and the image that was originally linked on the page. This should enable efficient multimodal learning, where the pipeline transparently solves the problem of matching an image with its associated text. To optimize the ratio of allocated CPU to GPU resources depending on the model, a profiling method based on the TensorFlow Profiler is provided.

REFERENCES

- [1] H. Yang, L. Liu, I. Milligan, N. Ruest, and J. Lin, “Scalable content-based analysis of images in web archives with tensorflow and the archives unleashed toolkit,” in *JCDL 2019*.
- [2] J. Bevendorff, M. Potthast, and B. Stein, “Fastwarc: Optimizing large-scale web archive analytics,” in *OSSYM 2021*.

Contact: <firstname>.<lastname>@uni-leipzig.de

This work was supported by the German Federal Ministry of Education and Research (BMBF, 01IS18026B) by funding the competence center for Big Data and AI “ScaDS.AI” Dresden/Leipzig.

¹ <https://github.com/chatnoir-eu/chatnoir-warc-dl>

² <https://www.iso.org/standard/68004.html>

³ <https://github.com/chatnoir-eu/chatnoir-resiliparse>