The Expression of Temporality in Historical Photographs

Alexandra Barancová Melvin Wevers Nanne van Noord

CANAL Lab, University of Amsterdam

Abstract

For this work we explore how temporal information is expressed in images by means of computer vision models, focusing specifically on historical photographs. Our study is centered around OpenCLIP, a vision-language model trained on extensive web-data, known for its capacity for zero-shot classification—a characteristic especially useful for cultural applications where training data is scarce. We scrutinize the sensitivity of such models to temporal cues present in images.

Our experiments employ the *De Boer Scene Detection* dataset, a collection of 39,866 gray-scale historical press photographs from 1950 to 1999. This dataset is a subset of the larger collection of the *De Boer* press agency, acquired in 2013 by the Dutch archive *Noord-Hollands Archief*. Noteworthy for its regional and national relevance, the agency chronicled a wide array of events ranging from local sports matches and construction projects to nationally significant occurrences like the 1953 North Sea Flood and the only two performances of The Beatles in the Netherlands.

Our image dating results show that zero-shot classification was relatively ineffective for precise image dating and exhibited a tendency for predicting dates in the past. However, we found that fine-tuning could mitigate this issue. In our exploration of zero-shot classification, as depicted in Figure 1, we discerned that colorization had a significant impact on dating results. Our in-depth analysis further uncovered that images containing people were dated more accurately, hinting at the existence of temporal markers.

This study underscores the potential of machine learning models like OpenCLIP for dating images and reiterates the importance of fine-tuning for accurate temporal analysis. Future research should explore the application of these findings to color photographs and diverse datasets.

Figure 1: Examples of original (left) and colorized (right) photographs from the *De Boer Scene Detection* dataset for which colorization had the largest impact on decreasing error for zero-shot classification. Both images showed a large prediction error in the gray-scale variant, albeit in different directions. Colorization decreased the error from -37 to 0 for the top image and from 37 to 2 for the bottom image.