

Optical music recognition (OMR) is a niche subfield of computer vision, where some labeled datasets exist, but there is an order of magnitude more unlabeled data available. Recent advances in the field happened largely thanks to the adoption of deep learning. However, such neural networks are trained using labeled data only. Semi-supervised learning is a set of techniques that aim to incorporate unlabeled data during training to produce more capable models. We have modified a state-of-the-art object detection architecture and designed a semi-supervised training scheme to utilize unlabeled data. These modifications have successfully allowed us to train the architecture in an unsupervised setting, and our semi-supervised experiments indicate improvements to training stability and reduced overfitting.