

We explore a dataset of almost half a million English song lyrics through three different processes — automatic evaluation, visualization, and generation. We create our own rhyme detector, using the EM algorithm with several improvements and adjustable parameters. This may, in some cases, replace human evaluators that cannot be used, for example, after each iteration of the lyrics generator to evaluate its improvement. By creating a web-page visualization of the results with interesting matrix rhyme highlighting, we make our evaluation accessible to the public. We discuss interesting genre differences discovered by applying our automatic evaluation on the entire dataset. Finally, we explore lyrics generation using state-of-the-art GPT-2.