POEM-KD: Persona-Oriented Efficient Model for Poetry Rewriting With Knowledge Distillation

Kim, Katie (School: Deerfield Academy)

Poetry rewriting requires more than fluent paraphrasing; it must preserve the source poem’s meaning while reproducing a target author’s persona, including diction, cadence, punctuation, and symbolic tone. Existing controllable text generation methods often retain surface meaning but fail to sustain a convincing poetic voice, while standard knowledge distillation improves efficiency at the cost of stylistic depth. To address this gap, we propose POEM-KD, a Persona-Oriented Efficient Model for Poetry Rewriting with Knowledge Distillation. POEM-KD uses a hierarchical Teacher–Assistant–Student framework, where a large Teacher transfers semantic and stylistic knowledge to an intermediate Assistant and then to a lightweight Student model. The Student is trained with a multi-objective loss combining cross-entropy supervision, knowledge distillation, semantic preservation, and persona fidelity. Training and evaluation use public-domain English poems from Wikisource and Project Gutenberg, alongside a separate test-only set of self-authored poems to assess out-of-distribution generalization. Results show that POEM-KD achieves a semantic similarity score of 0.91 and a persona fidelity score of 0.88, outperforming baseline systems while reducing inference latency to 64 ms through quantization, pruning, and adapter-based tuning. Qualitative analysis indicates that the model produces globally coherent persona transfer rather than isolated stylistic features. These findings suggest that hierarchical distillation with meaning- and persona-aware objectives enables efficient and scalable poetry rewriting across both in-distribution and out-of-distribution settings.

Awards Won: