

# LyricsBERT: Musical Recommendation on a Transformer-based Neural Embedding Backbone

Ronald Xu, Daniel Kim, Joey Zheng, Kevin Wen  
MIT 6.8611, Natural Language Processing, Fall 2023

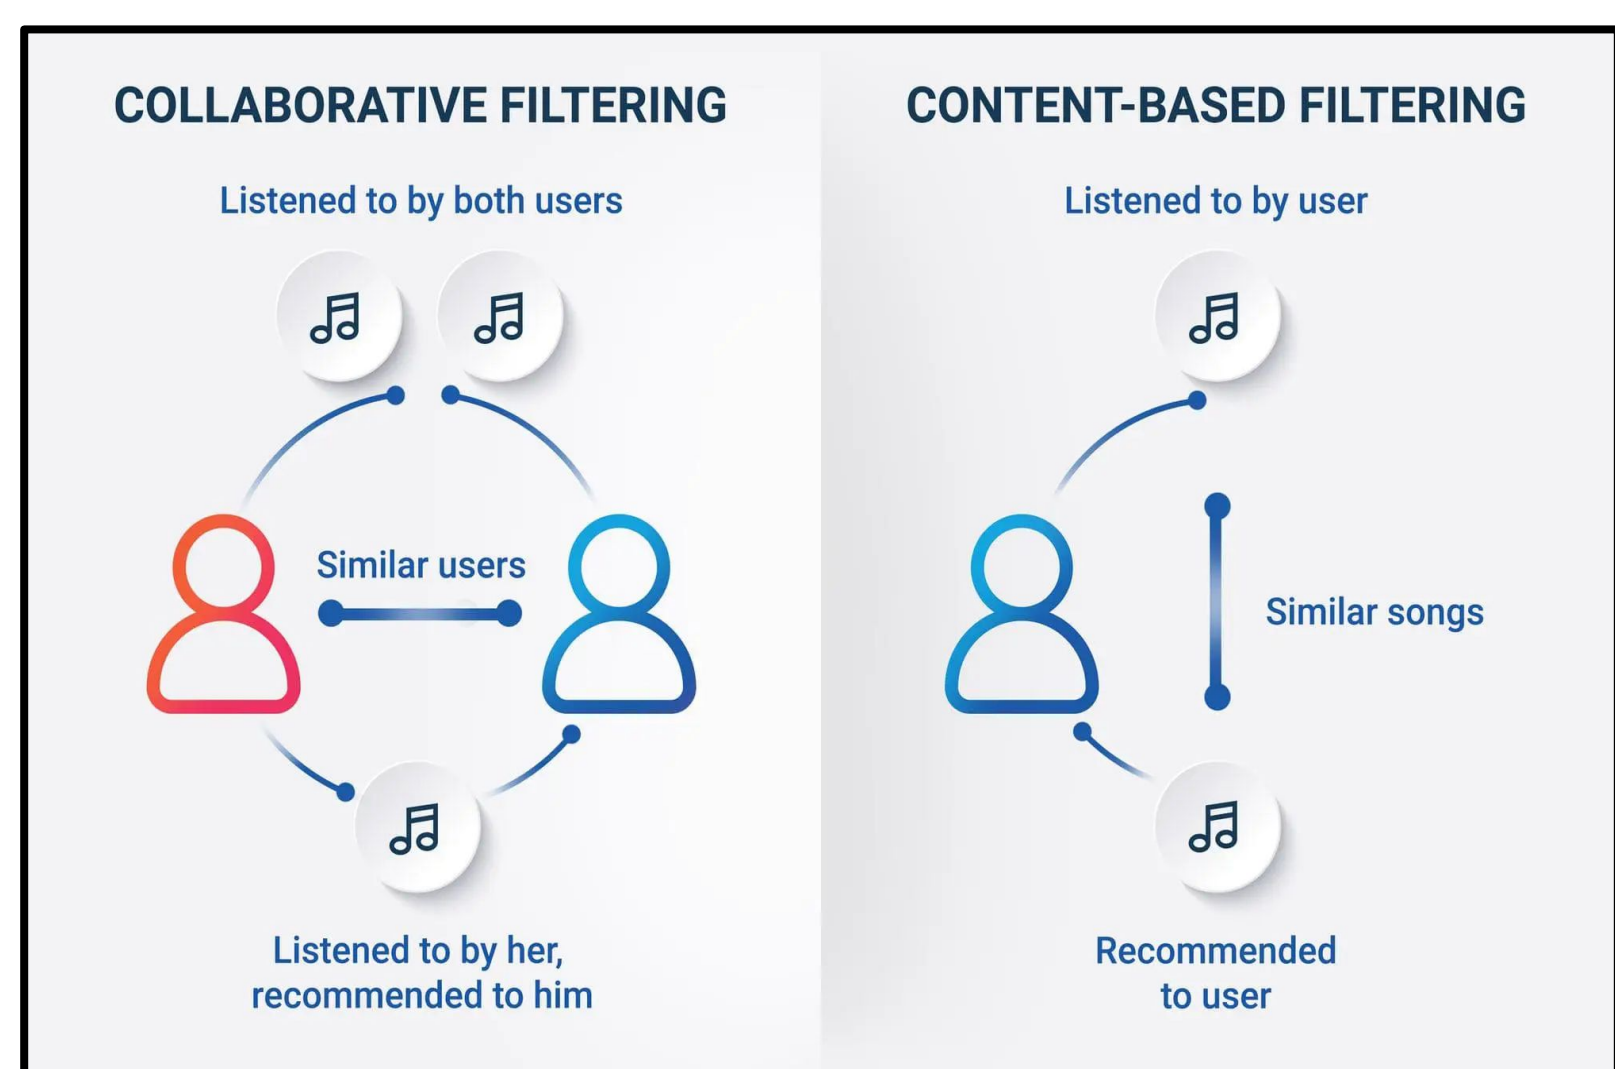


## Background & Motivation



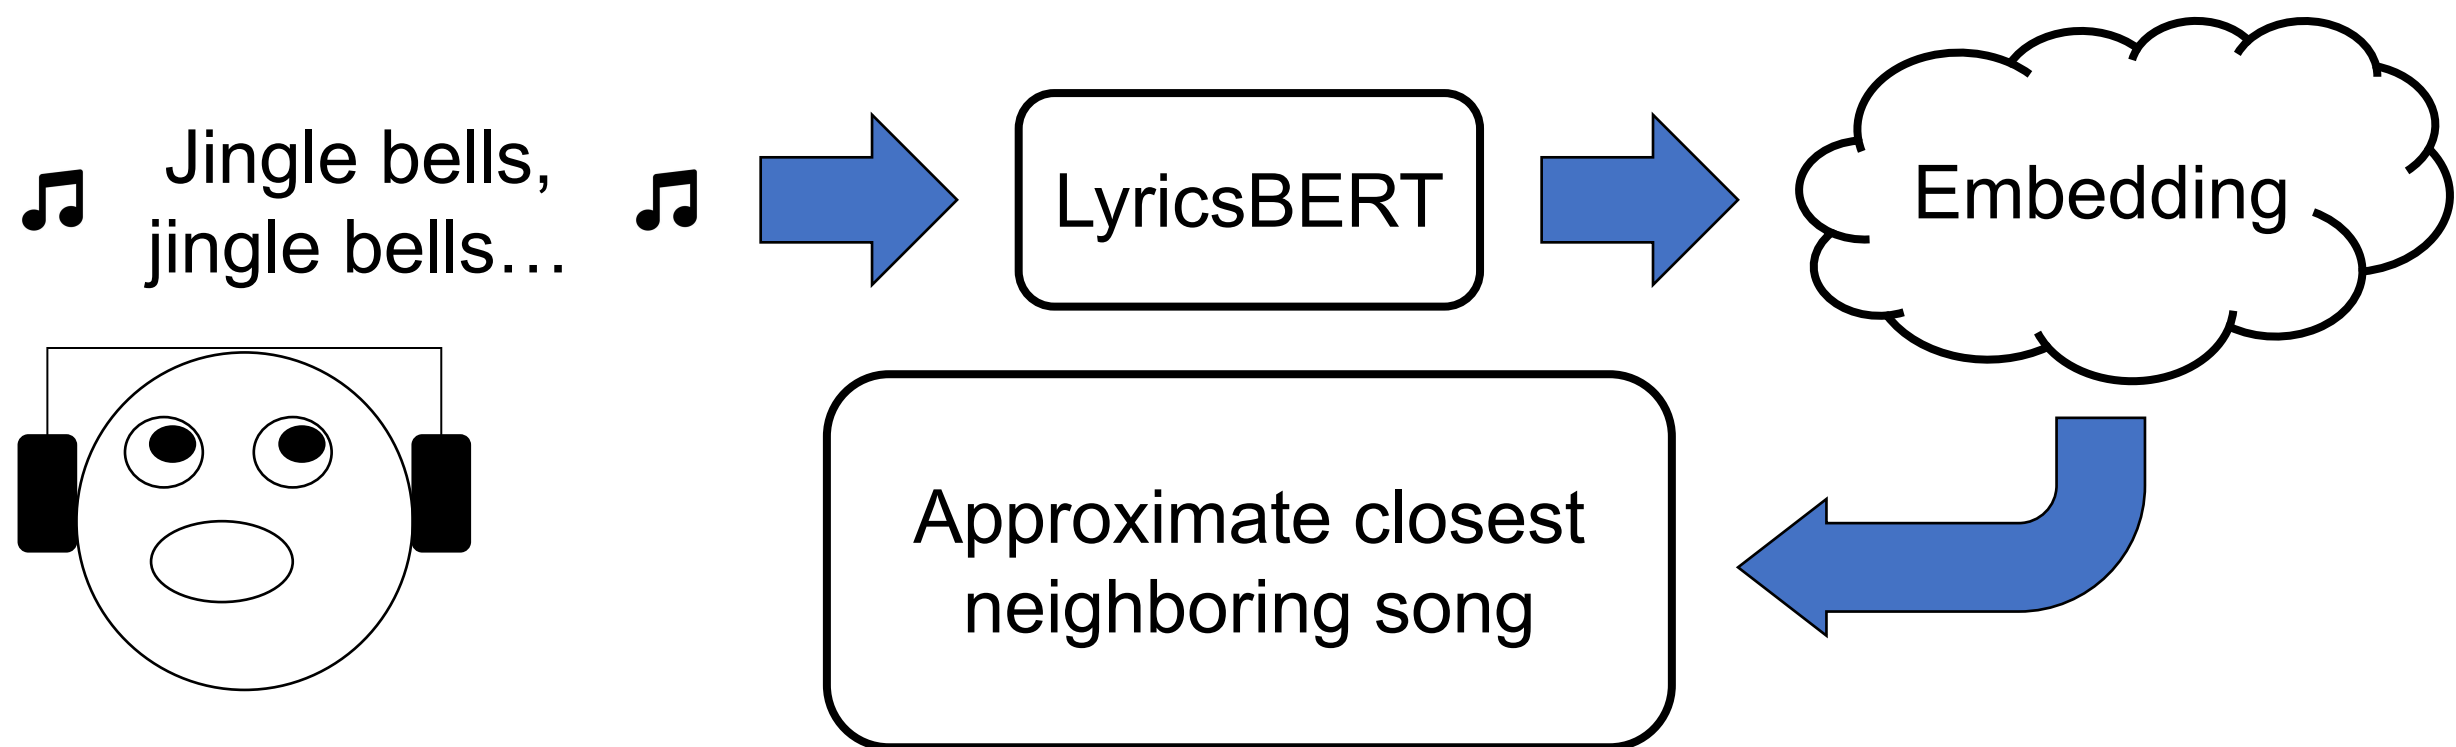
How can we create music recommendation system that **amplifies creativity**?

- Social-network:
- Ignores actual music
- Emotion-matching:
- Predicted emotions may be **too general** for recommendation

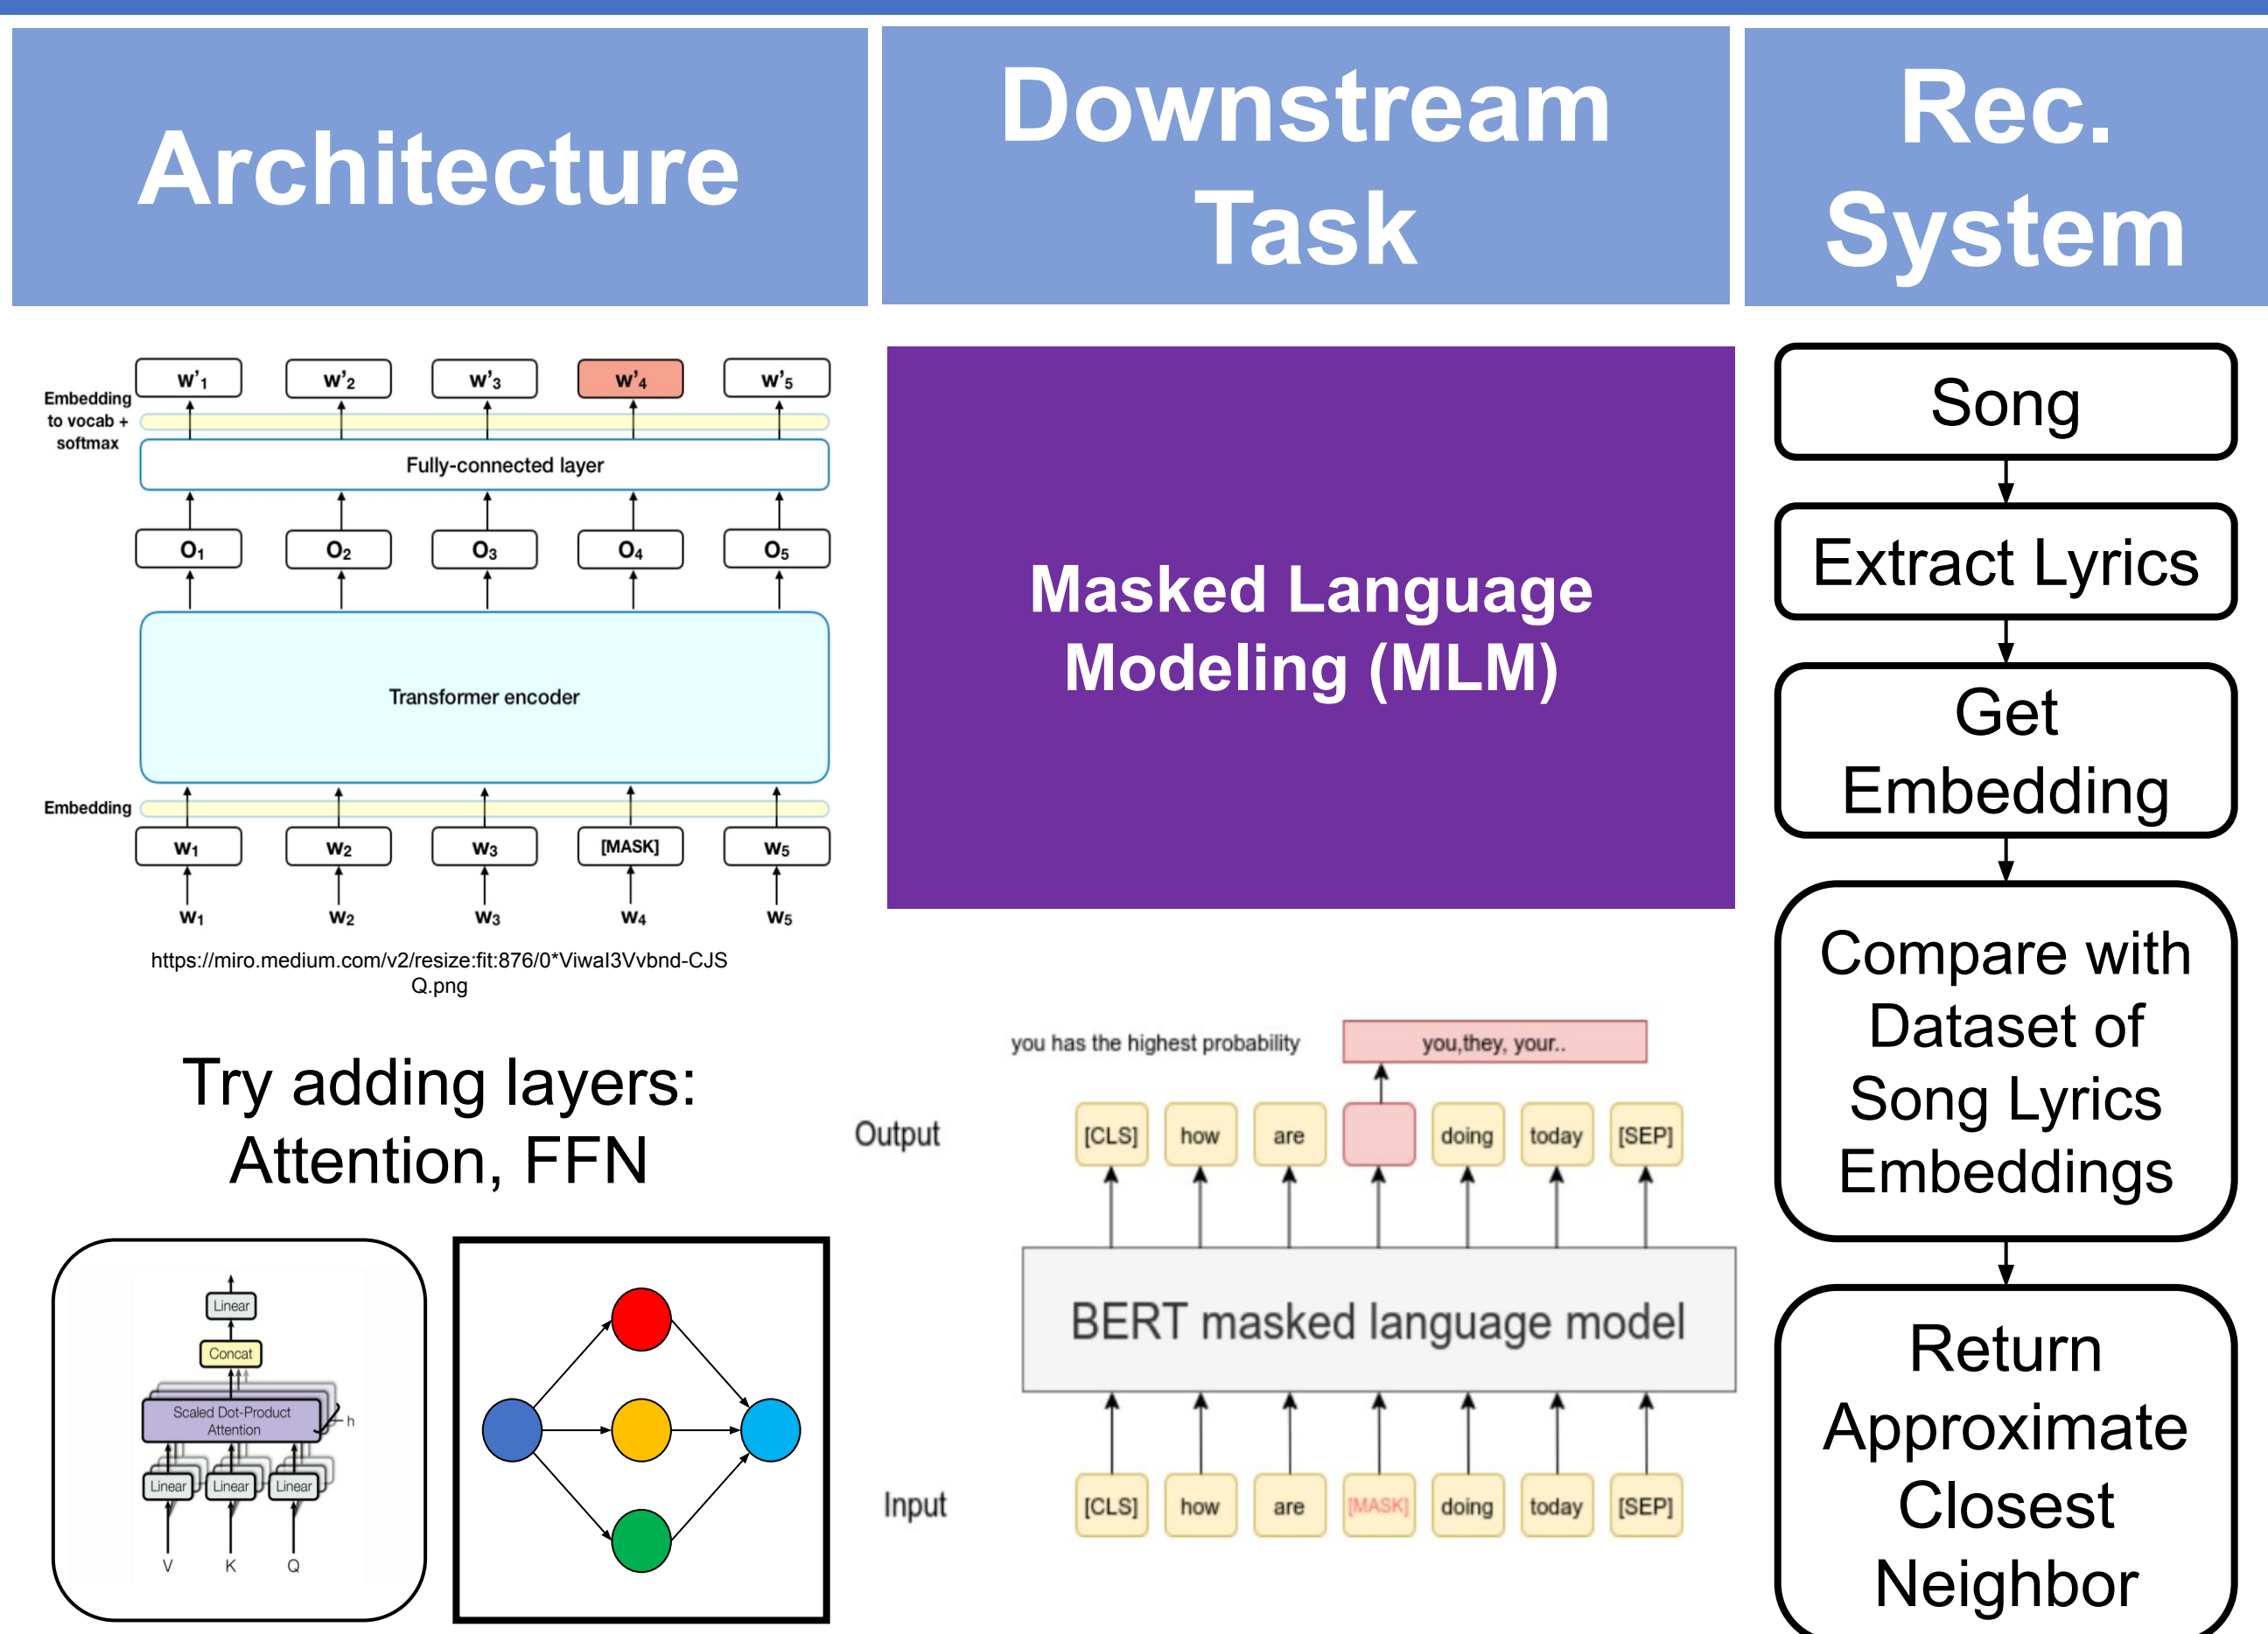


## Contribution

LyricsBERT: Unlock **lyrical meaning** with **specific** features

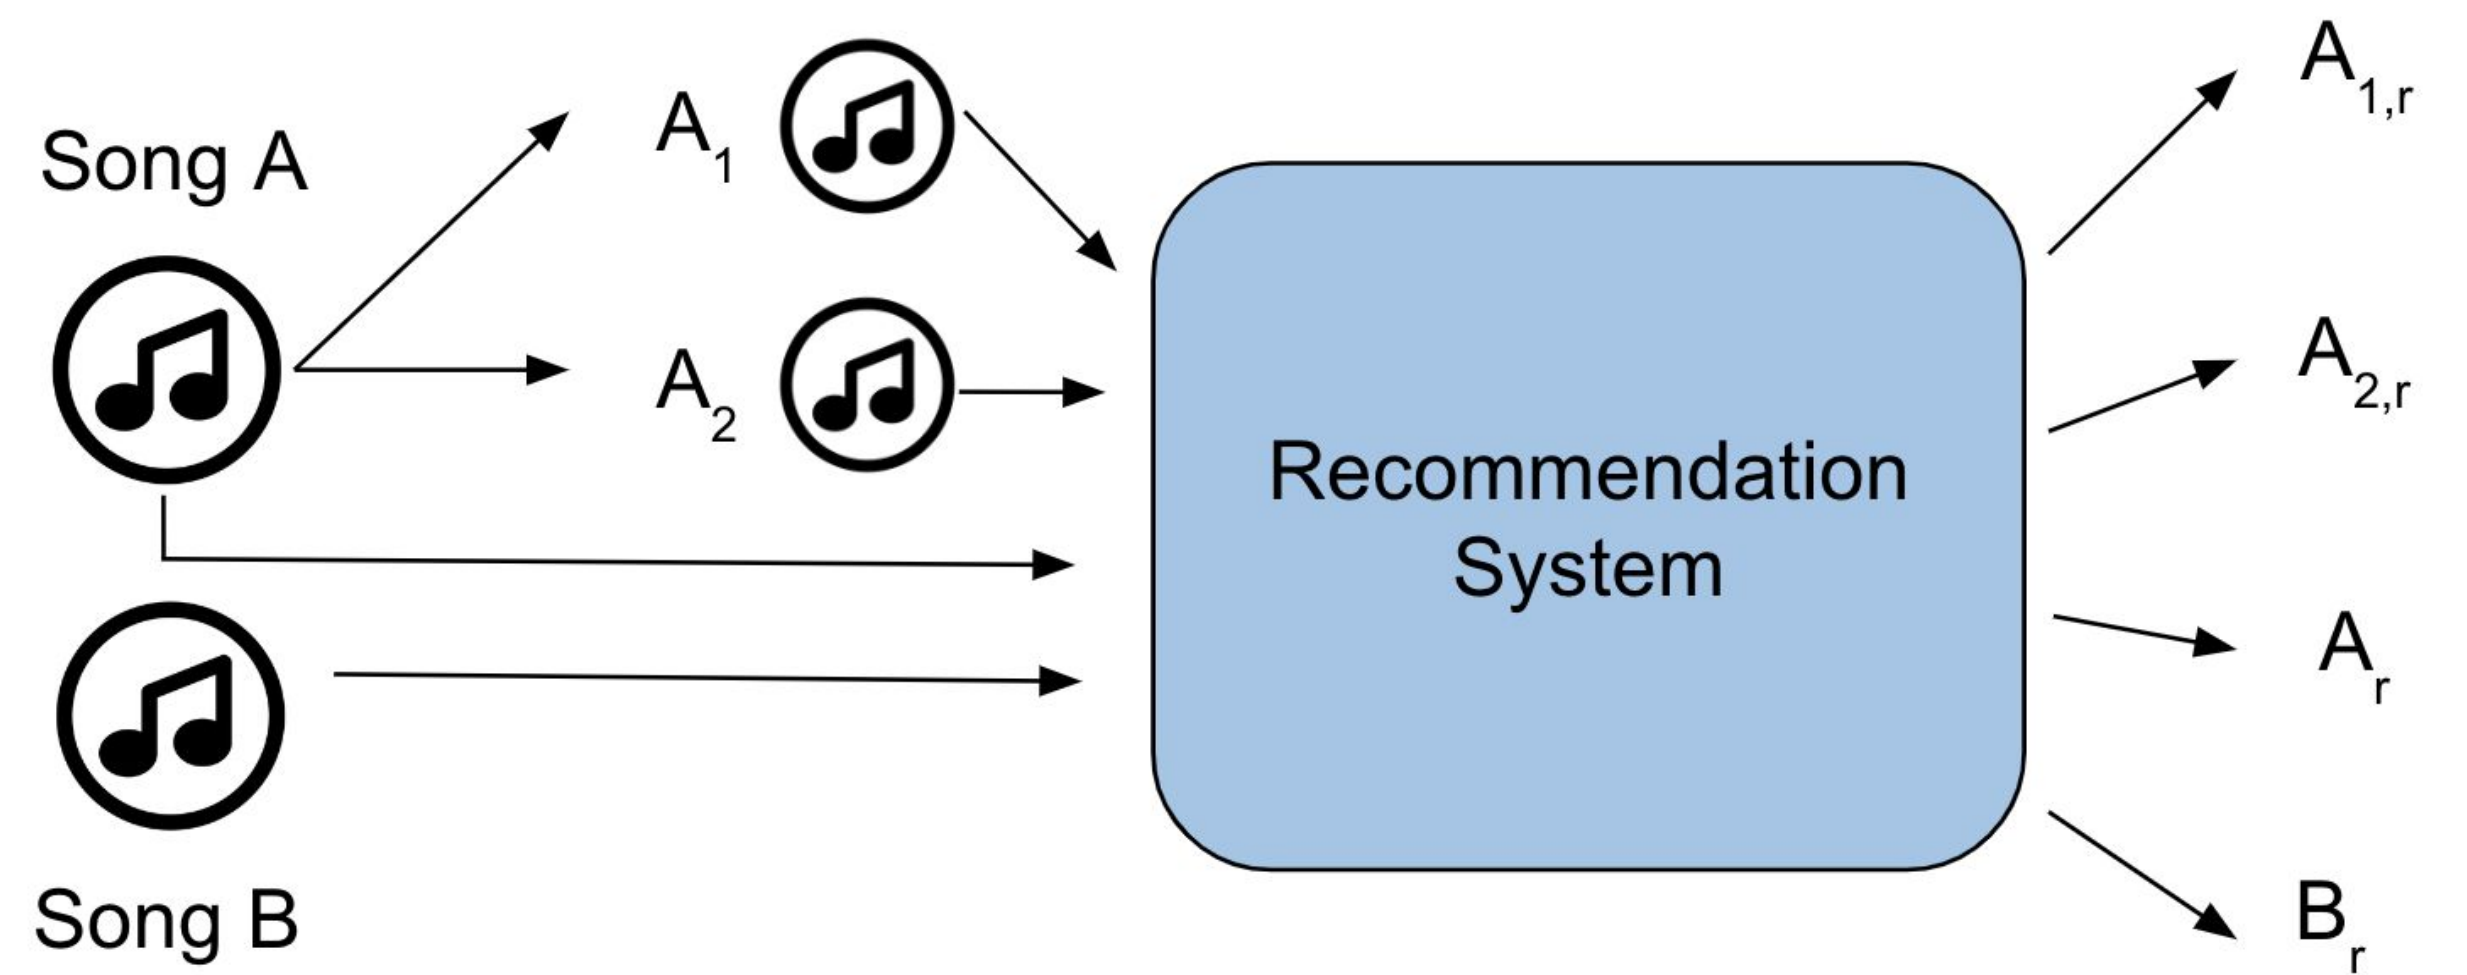


## Methods



## Evaluation


Recommendation system should output songs **close in embedding space** when input songs' **lyrics are similar**



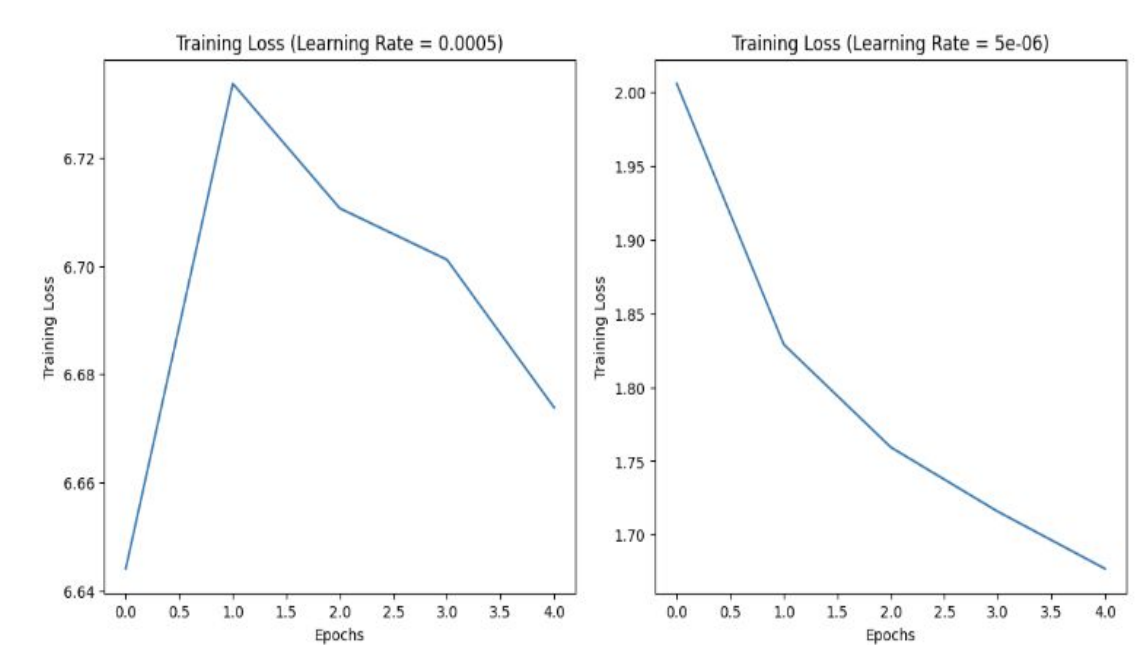
$$S = \|( |e(A_r) - e(B_r)| - |e(A_{1,r}) - e(A_{2,r})| )|_1$$

## Results & Discussion


Average Embedding Difference vs. Learning Rate



Original Architecture



Modified Architecture



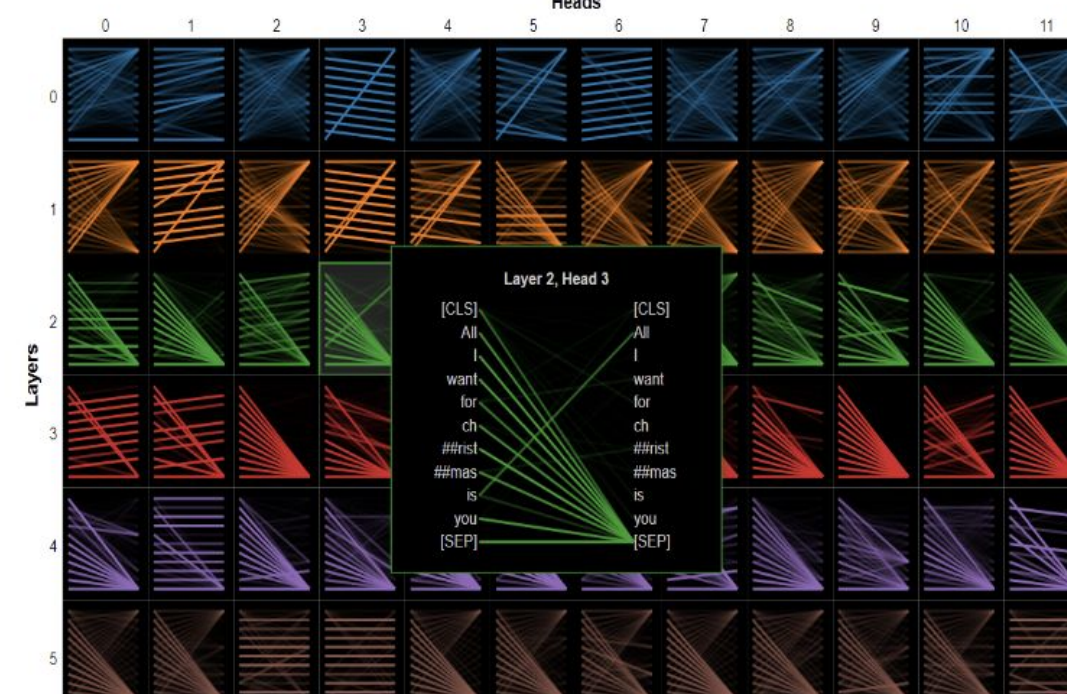
DistilBERT with added layers performs better

With a smaller learning rate, the original architecture learns better

With a larger learning rate, the modified architecture learns better

### Attention Visualization

Original Architecture



Modified Architecture

