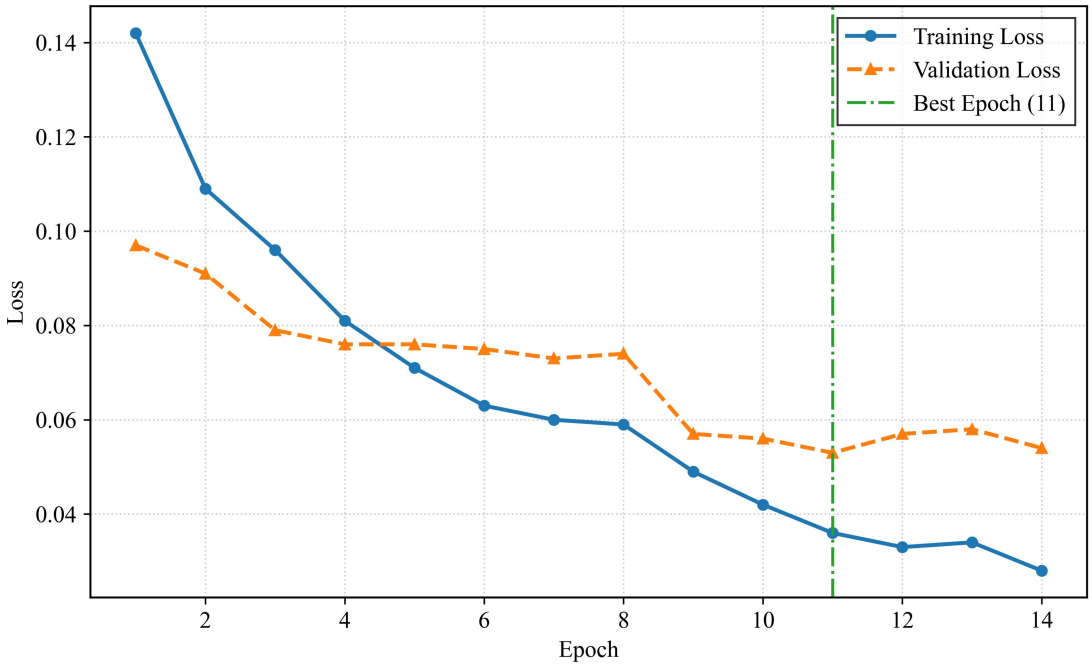


Supplementary Figure S4: Training and validation loss of the node scoring model. We report the mean squared error loss Lrank (Eq. (1) in the main paper) over epochs. Early stopping with patience 3 selects epoch 11 as the best checkpoint, and training terminates after no further improvement. Training loss is computed in training mode (with dropout/regularization), while validation loss is computed in evaluation mode.