

Additional file 1

Table 1: **Table S1. Training configuration for all evaluated embedding models.** Rank ($d = 300$), batch size ($b = 512$), and maximum epoch budget ($E = 750$, with early stopping on training loss stabilisation) were shared across all models. Optimizer, learning rate, loss function, and regularisation follow each architectural family’s original formulation.

Model	Framework	Optimizer (LR)	Loss / Regularisation
T-TransE	TA_TransE (PyTorch)	Adam (10^{-3})	Margin-based
TA-TransE	TA_TransE (PyTorch)	Adam (10^{-3})	Margin-based
TA-DistMult	TA_TransE (PyTorch)	Adam (10^{-3})	Softplus
ComplEx	OpenKE (PyTorch)	Adagrad (10^{-2})	Cross-entropy + N3
T-ComplEx	tkbc (PyTorch)	Adagrad (10^{-2})	Cross-entropy + N3 + L3
TNT-ComplEx	tkbc (PyTorch)	Adagrad (10^{-2})	Cross-entropy + N3 + L3
TeAST	Custom (PyTorch)	Adagrad (10^{-2})	Cross-entropy + N3 + L3
TCompoundE	Custom (PyTorch)	Adagrad (10^{-2})	Cross-entropy + N3 + L3