

Supplementary Material for “RAPPPID: Towards Generalisable Protein Interaction Prediction with AWD-LSTM Twin Networks”

Joseph Szymborski¹ and Amin Emad^{1*}

¹ Department of Electrical and Computer Engineering, McGill University, Montréal, QC, Canada

* Corresponding Author:

Amin Emad

755 McConnell Engineering Building

3480 University Street

Montréal, QC, Canada, H3A 0E9

Email: amin.emad@mcgill.ca

Supplemental Tables

Supplementary Table S1: Ranges of hyperparameters used for training of RAPPID. This table shows the range of hyperparameters and their increments considered when training RAPPID. The final values of hyperparameters were chosen based on the performance on a randomly selected subset of the training set (i.e., a validation set).

Hyperparameter	Number of LSTM Layers	Embedding Dropout Rate	LSTM Dropconnect Rate	Classifier Dropout Rate	Learning Rate
Range	2-3	0.1-0.4	0.1-0.4	0.1-0.4	$10^{-3} - 10^{-2}$
Increment	1	0.1	0.1	0.1	0.009

Supplementary Table S2: The final values of hyperparameters selected for the analysis reported in Table 1 of the manuscript for different datasets. The hyperparameters are selected based on the performance of a validation set (separate from the testing set).

Experiment	Number of LSTM Layers	Embedding Dropout Rate	LSTM Dropconnect Rate	Classifier Dropout Rate	Learning Rate
C1	3	0.1	0.1	0.1	10^{-2}
C2	2	0.3	0.3	0.2	10^{-2}
C3	2	0.3	0.3	0.2	10^{-2}

Supplementary Table S2: The final values of hyperparameters selected for the analysis reported in Table 2 on dataset C3. The hyperparameters are selected based on the performance of a validation set (separate from the testing set).

Experiment	Number of LSTM Layers	Embedding Dropout Rate	LSTM Dropconnect Rate	Classifier Dropout Rate	Learning Rate
Repeatability 1	2	0.3	0.2	0.2	10^{-2}
Repeatability 2	2	0.3	0.3	0.2	10^{-2}
Repeatability 3	2	0.2	0.3	0.2	10^{-2}

Supplementary Figures

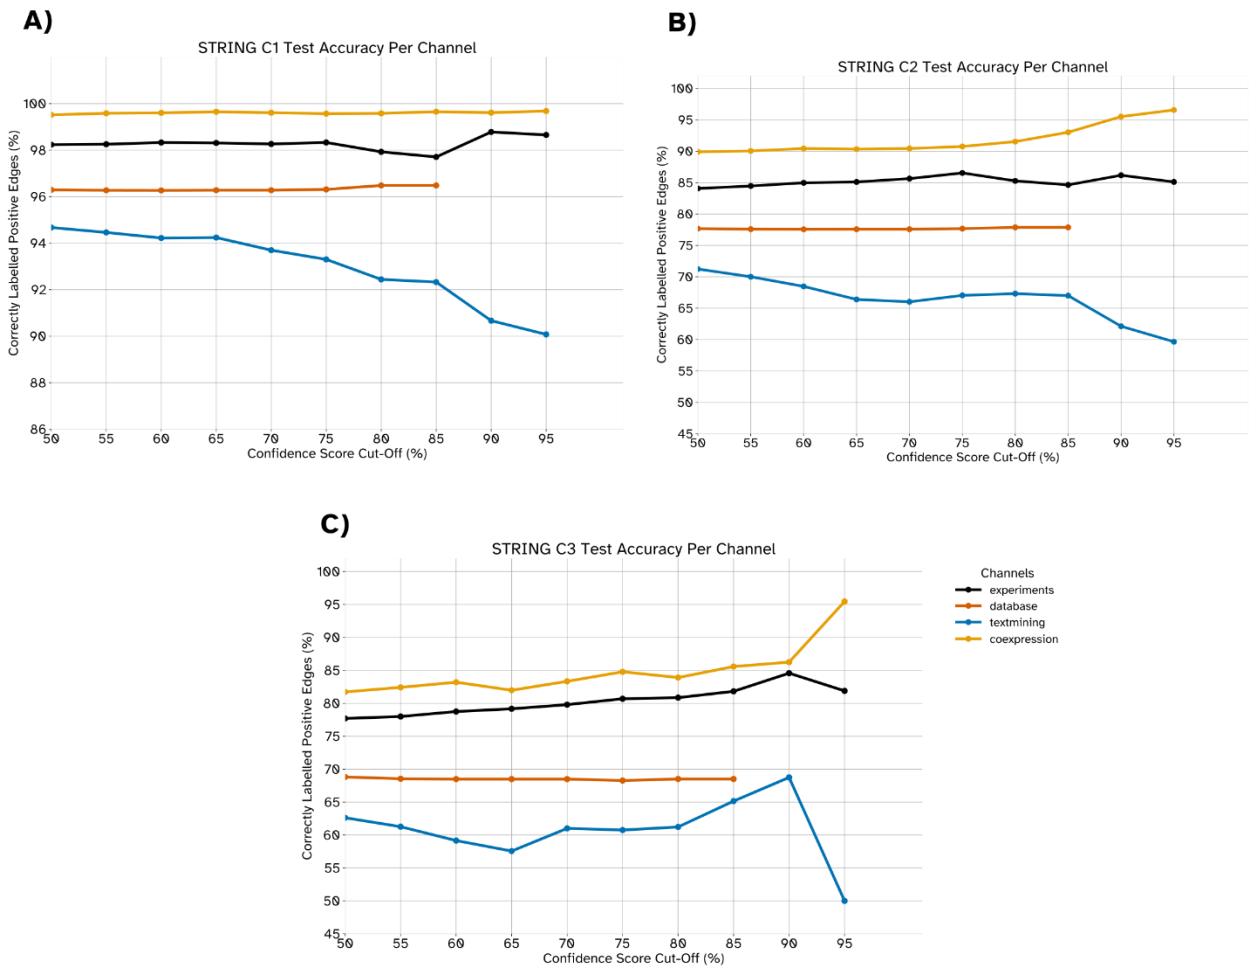


Figure S1: Accuracy of positive edges across edge confidence stratified by STRING channels. The percentage of correctly labelled positive edges are plotted for each major STRING channel in the C1 (A), C2 (B), and C3 (C) datasets. The x-axis denotes the channel edge confidence cut-off score for each curve's respective channel.

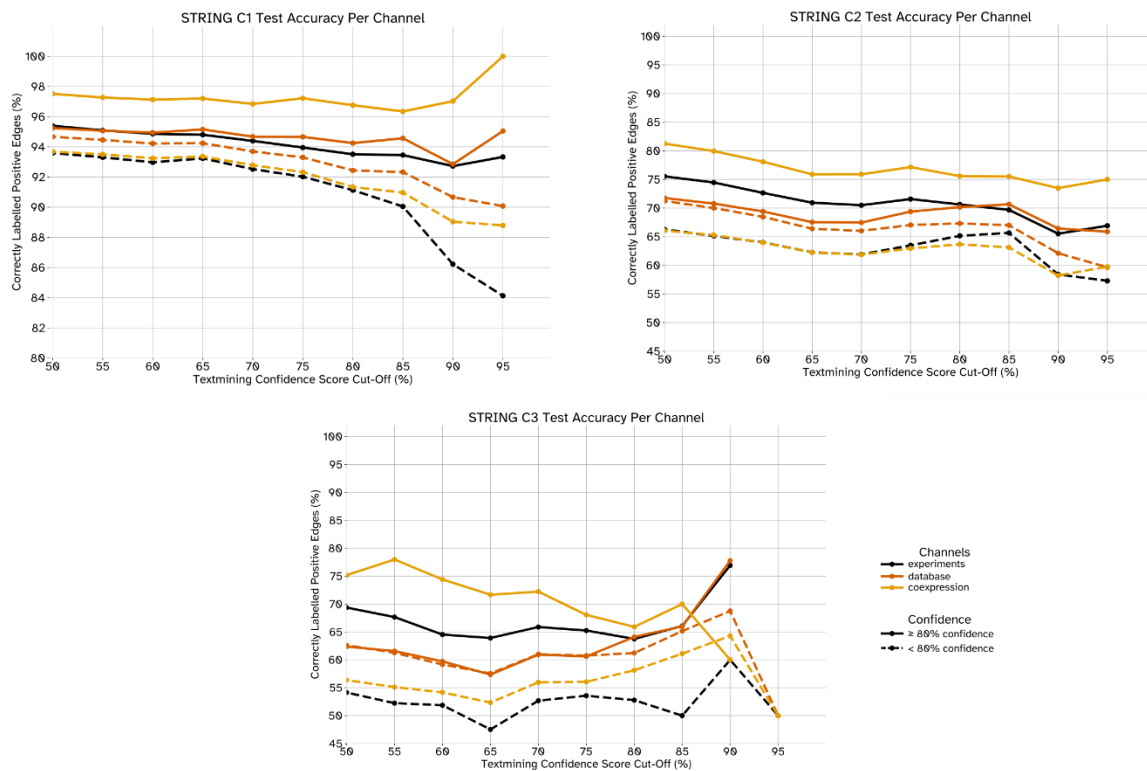


Figure S2: Accuracy of positive edges across edge confidence stratified by STRING channels. The percentage of correctly labelled positive edges are plotted for each major STRING channel, excluding text-mining. Edges are excluded according to the text-mining confidence threshold (x-axis). The solid curves include edges which have a channel confidence $\geq 80\%$ for the channel indicated by the curve's colour. Dashed curves conversely include edges whose channel confidence is $< 80\%$ for the channel indicated by the curve's colour. Data shown reflects the C1 model/dataset in panel (A), the C2 model/dataset in panel (B), and the C3 model/dataset in panel (C).




Figure S3: Accuracy of positive edges as a function of similarity between testing and training proteins in C3. The similarity between testing and training proteins was measured using their percent identity as computed by NCBI’s PSI-BLAST software. The highest percent identity between any training protein and a testing protein in a testing edge was considered to be that testing edge’s “maximum percent identity”. The percentage of accurately labelled positive edges (black curve, left y-axis) is reported for edges with maximum percent identities lower than the threshold reported on the x-axis. The proportion of testing edges for each threshold values is reported by the dashed blue curve and the right y-axis.