

## 640 A Appendix

### 641 A.1 Experiment Configurations

642 To obtain our REFuSe model, we initialized the neural net with an embedding size of 8, a window  
643 size of 8, a stride size of 8, and an output size of 128. We trained the model for 30 epochs over the  
644 Assemblage training dataset, with the model seeing 10M functions per epoch. Functions were divided  
645 into batches of 600; for each batch, 300 unique labels were randomly chosen from the training dataset,  
646 and then two functions were randomly selected with each label. We used a learning rate of 0.005, and  
647 the Adam optimizer with gradients clipped to  $[-1, 1]$ . Per the literature [18; 14], we used  $\alpha = 0.2$  as  
648 the margin for our triplet loss. REFuSe was trained on three Tesla M40s on an internal cluster, and  
649 took 4.5 days to train.

650 To evaluate the GNN, we used the model checkpoint published by [33] as part of their survey<sup>1</sup>. To  
651 evaluate jTrans, we similarly used the fine-tuned model made available on the authors' Github page<sup>2</sup>.

### 652 A.2 Evaluation Procedures

653 We chose to use mean reciprocal rank (MRR) to measure how models performed on our benchmark.  
654 MRR, a popular metric in information retrieval, is used to assess systems which take in queries and  
655 return a list of possible responses ordered by likelihood of correctness. Letting  $q$  be a query,  $L$  be the  
656 1-indexed list returned by  $q$ , and  $c$  be a correctness function, where  $c(L[i]) = 1$  if  $L[i]$  is a correct  
657 response to  $q$  and 0 otherwise,  $q$  is said to have rank  $r$  if the first correct answer in  $L$  appears at  
658 position  $r$ . That is,  $q$  has rank  $r$  if and only if  $1 \leq r \leq \text{len}(L)$ ,  $c(L[r]) = 1$ , and  $c(L[i]) = 0$  for all  
659  $1 \leq i < r$ . The reciprocal rank of  $q$  is defined to be  $\frac{1}{r}$ , and the mean reciprocal rank is the average of  
660 the reciprocal ranks for every query  $q \in Q$ . The upper bound on MRR is 1.0 (a correct answer is  
661 always in the first position in  $L$ ), whereas the lower bound on MRR is 0.

662 In the context of BFS,  $q$  is a query function and  $L$  is a list of neighboring functions (embeddings),  
663 ordered from nearest to farthest. Due to the large size of our datasets, we used the Hierarchical  
664 Navigable Small Worlds [32] approximate nearest neighbor index from Faiss [12] to compute the 30  
665 nearest neighbors to each query function. When no match was found within the first 30 neighbors,  
666 we assigned that query an upper bound reciprocal rank of  $\frac{1}{31}$  and a lower bound reciprocal rank of 0.

667 In Section 5.1 we reported the lower and upper bound MRR for the experiments that used our  
668 evaluation code. For benchmarking experiments that utilized open-source code from other authors,  
669 we reported a single MRR value, keeping with their practice. In particular, when conducting  
670 experiments with jTrans, we chose to use the evaluation code published by its authors, as integrating  
671 our own code into their codebase was not straightforward. jTrans supports evaluation over multiple  
672 pool sizes; in Section 5.1 we report results for pool size 10,000, as a larger pool size more closely  
673 mimics the evaluation methods of the other models. (In our evaluation, the pool is the entire dataset,  
674 but we are not limited to having only one function matching the query function in each pool.)

## 675 B GNN Common Libraries Details

676 In our results we stated the significant drop in the GNN's performance on the Common Libraires  
677 corpus is due to its inability to handle the variety of functions and function sizes in each application.  
678 This is important to verify as the actual cause, as the asperity in the project sizes could easily dominate  
679 the results and make it unclear which method actually performs best.

---

<sup>1</sup>This model is available at the following link: [https://github.com/Cisco-Talos/binary\\_function\\_similarity/tree/main/Models/GGSNN-GMN/NeuralNetwork/model\\_checkpoint\\_GGSNN\\_pair](https://github.com/Cisco-Talos/binary_function_similarity/tree/main/Models/GGSNN-GMN/NeuralNetwork/model_checkpoint_GGSNN_pair).

<sup>2</sup>This model can be downloaded from <https://github.com/vul337/jTrans/>.

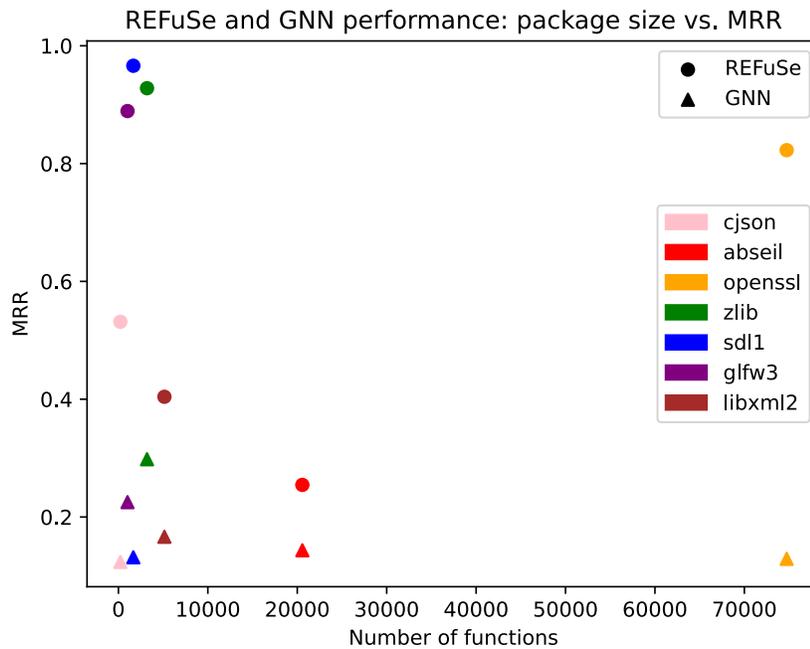


Figure 2: REFuSe and GNN per-package performance. The circles correspond to REFuSe results, while the triangles correspond to the GNN.

680 We perform this validation in [Figure 2](#), where it can be seen that REFuSe dominates the GNN in  
 681 performance for each library. Though there are too few libraries to make a definitive conclusion,  
 682 REFuSe seems to be unfazed by the number of functions in terms of final MRR performance. Yet,  
 683 the GNN has low performance in all cases and decreases with the number of functions.