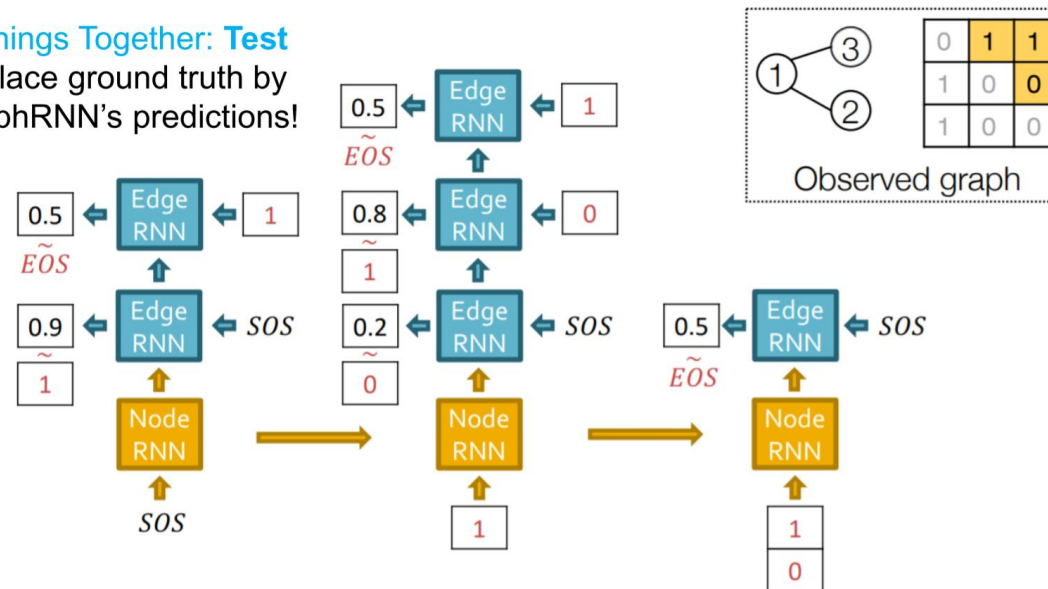


Summary Questions of the lecture

- What is the meaning of the output of each RNN cell in **GraphRNN**?
- The node-RNN at time-step t outputs the initial state of the edge-RNN. The edge-RNN sequentially outputs the probability that the current node t is connected to each existing node. The edge connectivity (1 for connected, 0 for not connected) is determined by sampling from the probability outputted by the edge-RNN. The edge-RNN stops when it outputs the end token EOS.

Put Things Together: Test

- Replace ground truth by GraphRNN's predictions!

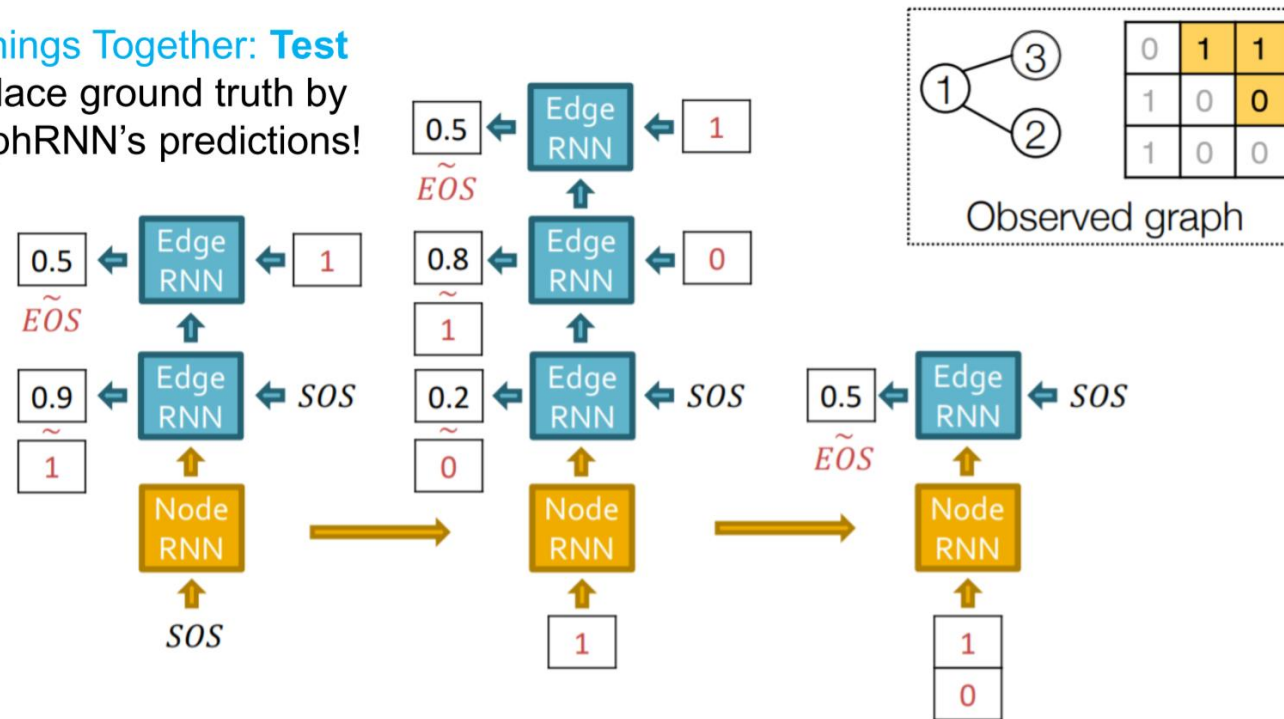


Summary Questions of the lecture

- How can we obtain the input of each RNN cell in **GraphRNN**?
- The first input of any RNN is the start token *SOS*. After that, the edge-RNN receives the binary output of the previous cell, and the node-RNN receives all of the outputs of the previous edge-RNN sequence as a vector.

Put Things Together: Test

- Replace ground truth by GraphRNN's predictions!



Summary Questions of the lecture

- Explain the training method of **GraphRNN** in the view point of loss and training path.
- During training, both RNNs receive the ground truth as input, regardless of the output of the previous cell. Output probabilities of edge-RNNs are learned in the supervised manner with the binary cross-entropy loss. Since RNNs weights are shared, the gradients w.r.t. the weights are accumulated across time-steps.

$$L = - \sum_i [y_i^* \log y_i + (1 - y_i^*) \log(1 - y_i)]$$

Put Things Together: Training

- Backprop:** All gradients are accumulated across steps

