

What?

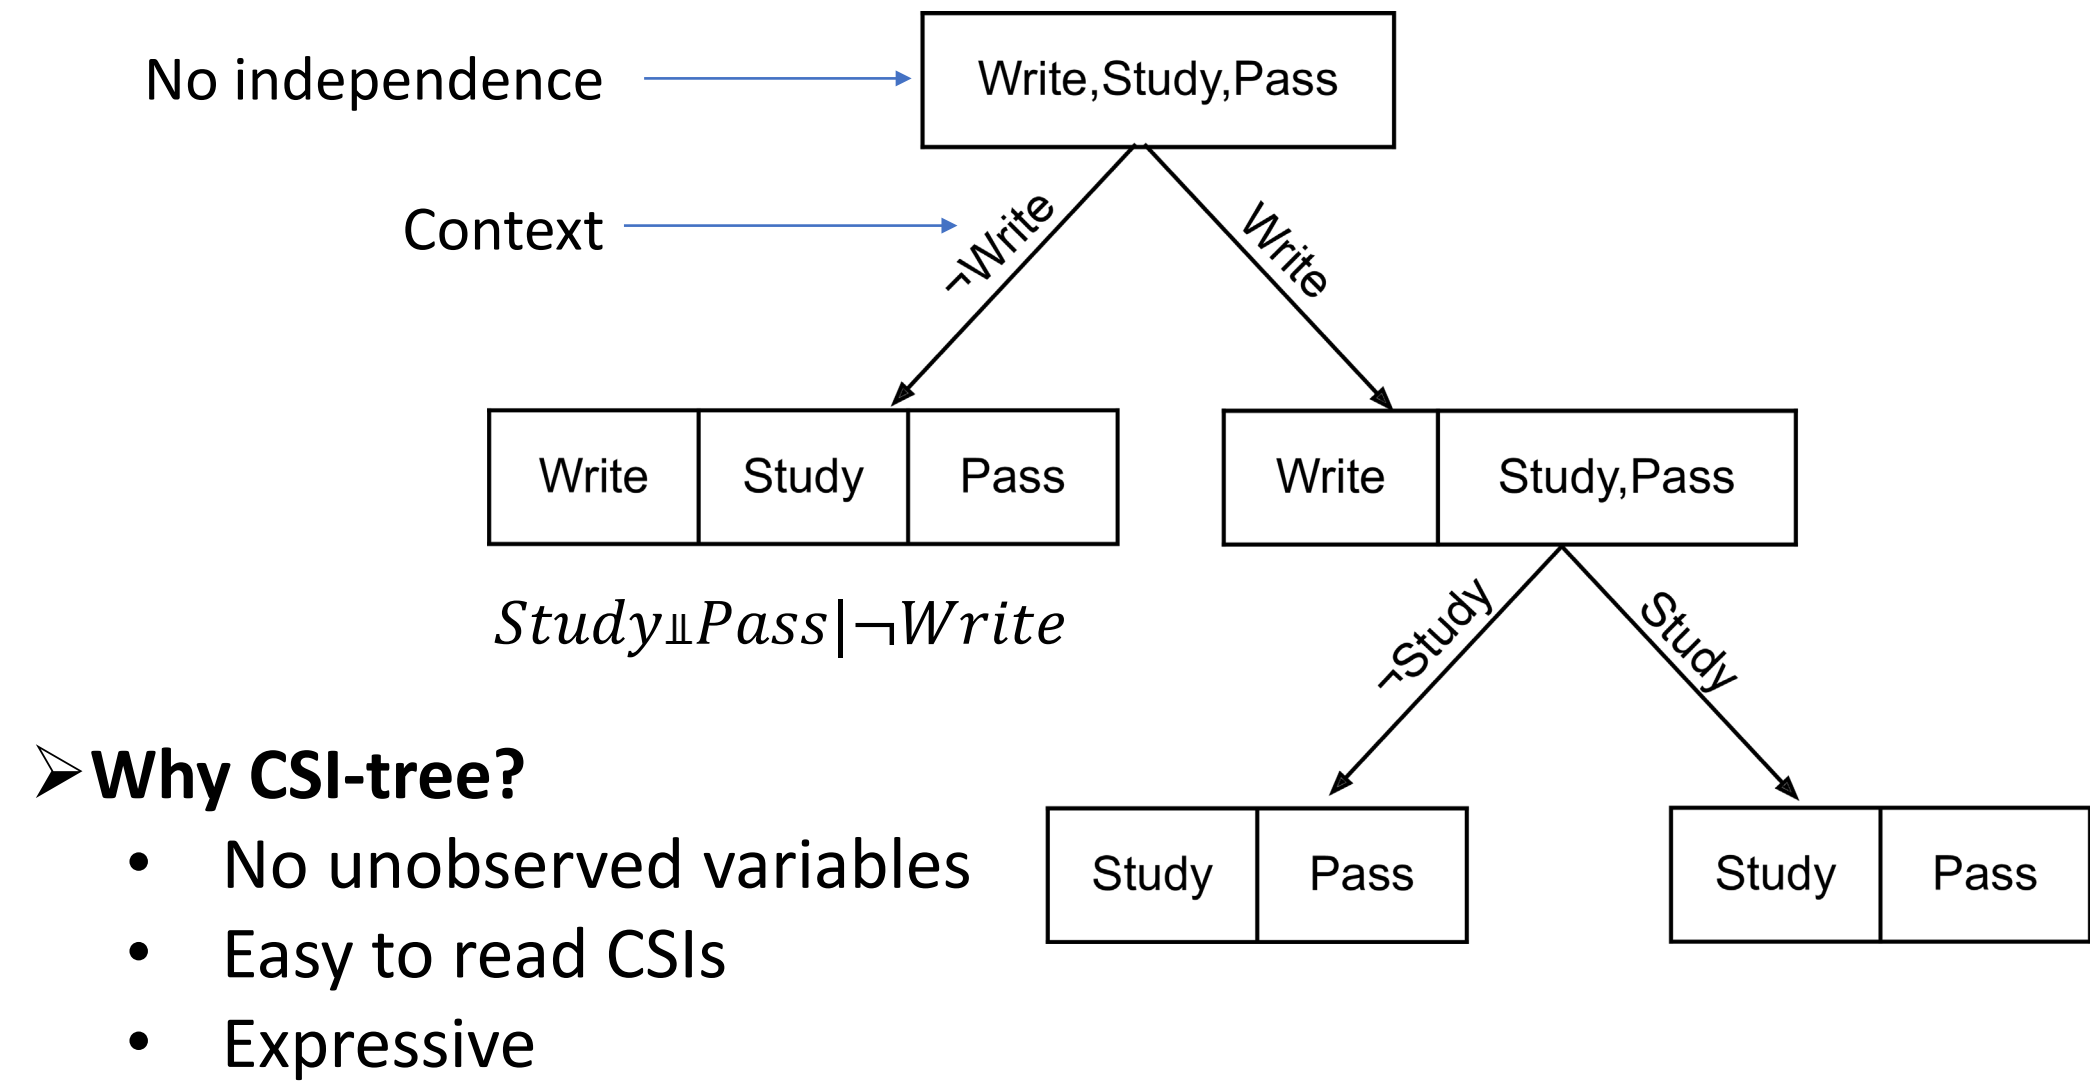
How?

Results

Motivation

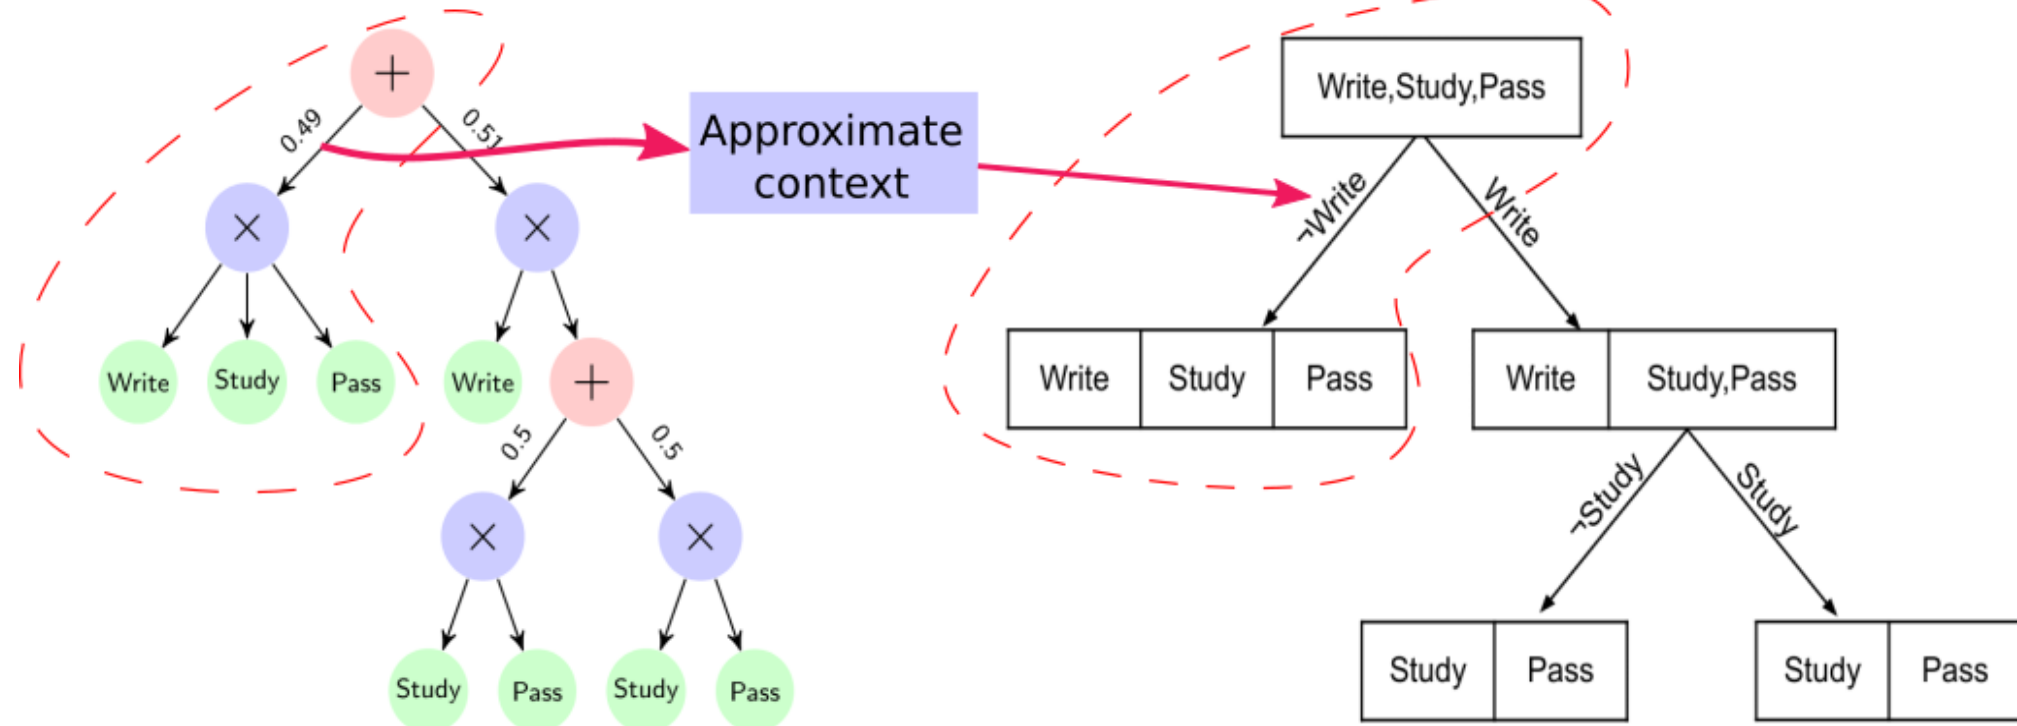
- Sum-Product Networks (SPNs) are inherently uninterpretable – internal nodes do not correspond to any feature
- Context-specific independences (CSIs) – Conditional independences that hold under certain instantiations of conditioned variables
- SPNs encode CSIs – product nodes capture independences
- Can convert SPNs into an interpretable representation by leveraging the CSIs encoded by the SPN?

CSI-tree

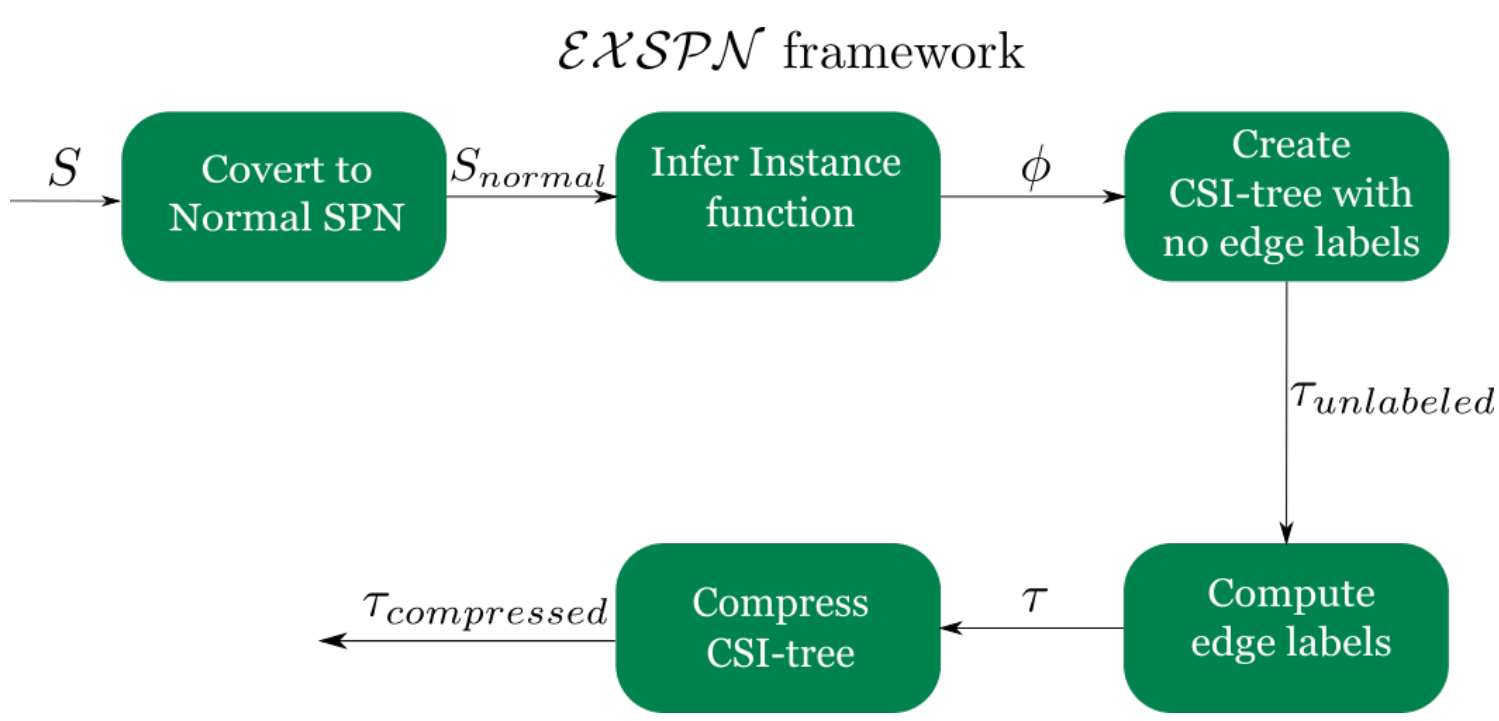


The ExSPN framework

Given: SPN $S = (G, \psi, w, \theta)$, data \mathcal{D}
To Do: Extract CSI-tree $\tau = (G_{csi}, \psi_{csi}, \chi, \zeta)$



- ExSPN:
- outputs CSI-tree that provably recovers the structure of the original normal SPN
 - is robust to noise – function approximator used to approximate the context is a generalizable discriminator
 - extracts CSI-tree that is linear in the size of a normal SPN



```

Algorithm 1: ExSPN
input :  $\mathcal{D}, \mathbf{X}, S = (G = (N, E), \psi, w, \theta), \lambda$ 
output : CSI-tree  $\tau = (G_{csi}, \psi_{csi}, \chi, \zeta)$ 
1 Convert  $S$  to a normal-spn  $S_{normal}$ 
2 Infer  $\phi$  using alg. 2
3 Initialize  $G_{csi} = G_{normal}, \psi_{csi} = \psi_{normal}, \tau = (G_{csi}, \psi_{csi}, \chi, \zeta)$ 
4  $st = [G_{csi}.root]$ 
5 while  $st$  is not empty do
6    $N_{current} = st.pop()$ 
7   if  $N_{current}$  is not the root node then
8     if  $N_{current}$  is a leaf node then
9       Add  $\psi(N_{current})$  to  $\chi(pa(N_{current}))$ 
10    end
11   if  $N_{current}$  is a sum node then
12     Delete  $N_{current}$  from  $G_{csi}$  Add
13      $\psi(N_{current})$  to  $\chi(pa(N_{current}))$ 
14     Connect parents and children of  $N_{current}$ 
15     Duplicate children with multiple parents and ensure tree-structure
16   end
17   Add  $ch(N_{current})$  to  $st$ 
18 end
19  $\tau = ComputeLabels(\tau, \mathcal{D}, \phi, \lambda)$ 
20 return  $\tau$ 
    
```

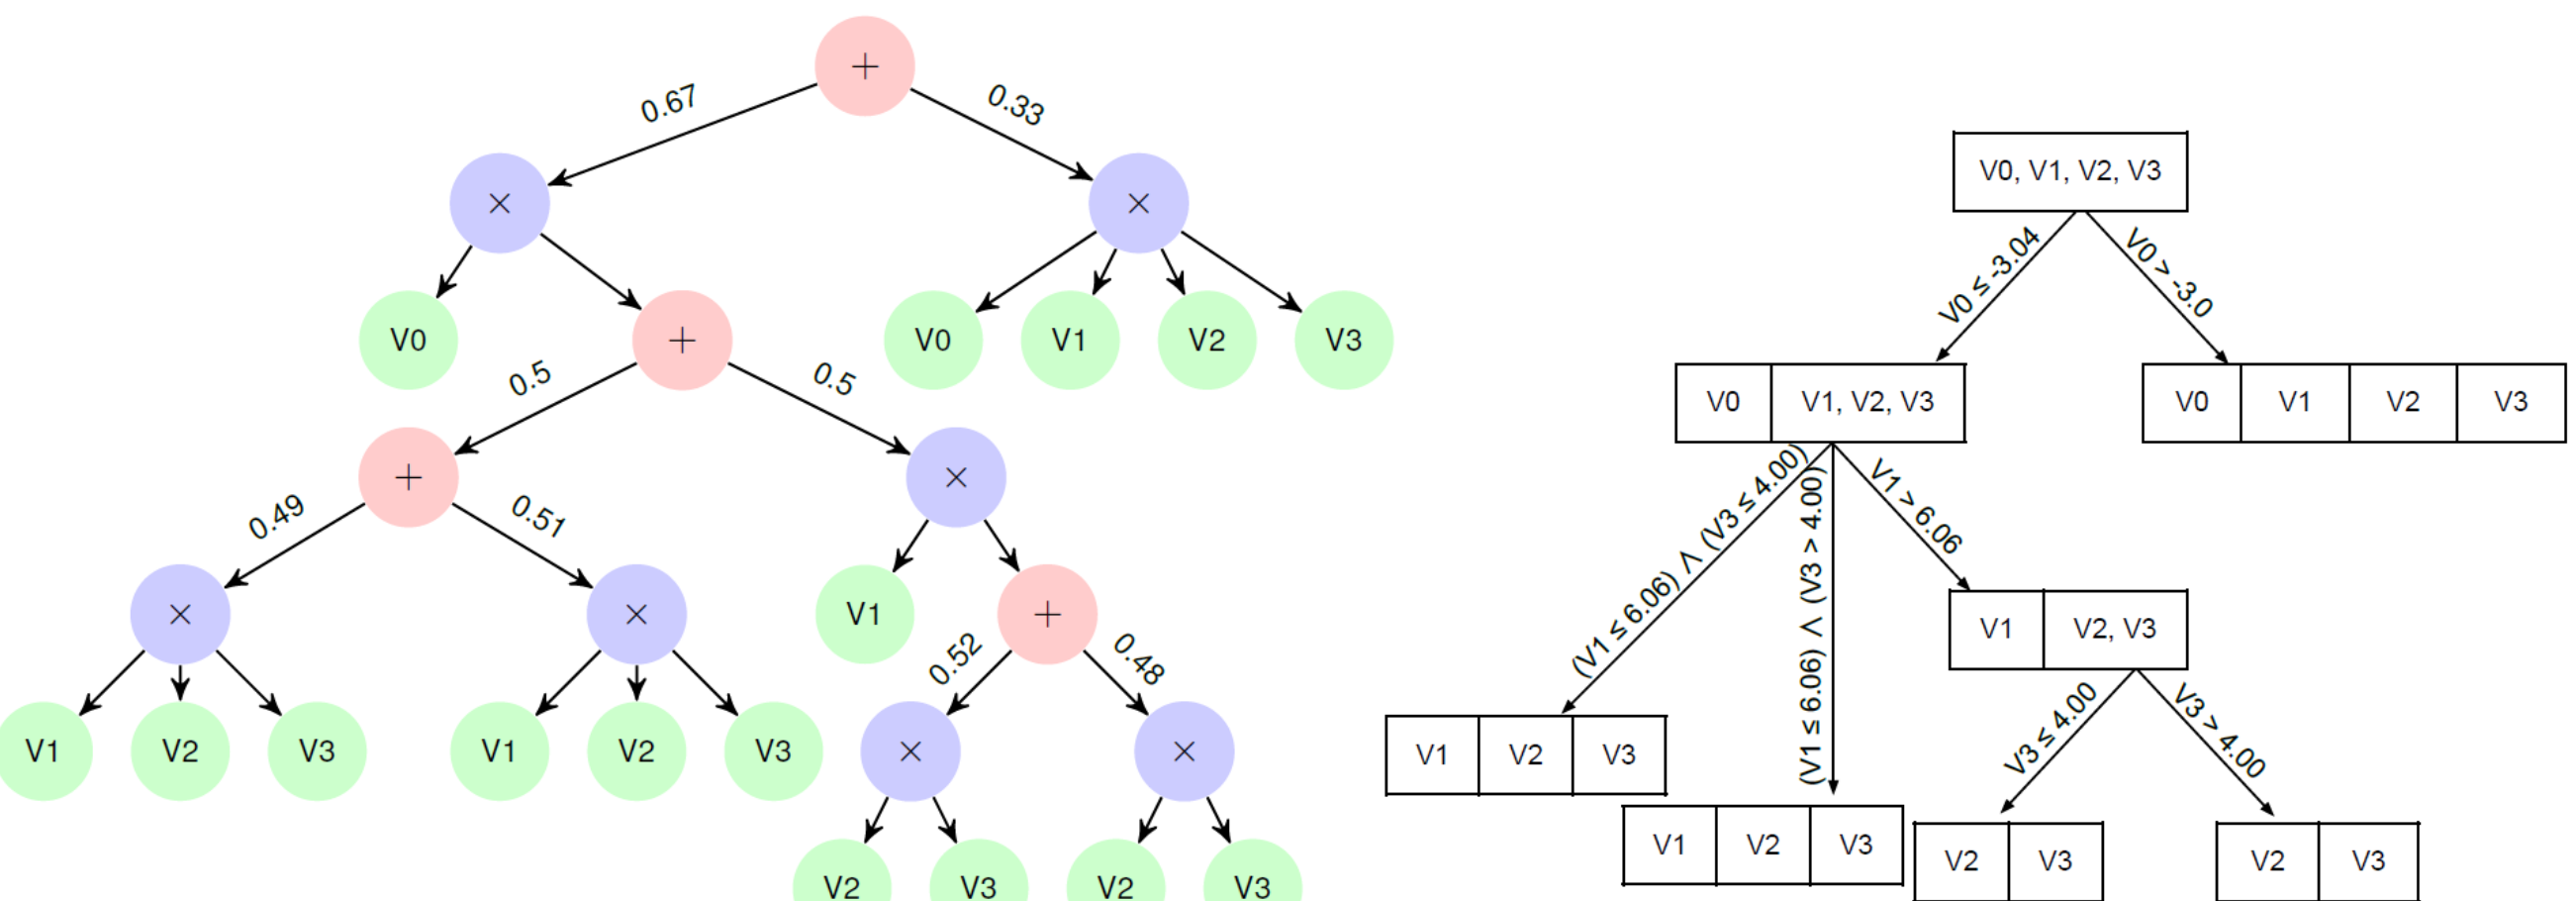
```

Algorithm 2: InferInstanceFunction
input :  $\mathcal{D}, \mathbf{X}$ , normal SPN  $S$ 
output : Instance function  $\phi$ 
1 for Each instance  $D_j \in \mathcal{D}$  do
2   Perform upward pass for  $\mathcal{D}$  and compute  $S_j^{max}(D_j)$  for each node  $N$ 
3   Perform a downward pass, appending  $D_j$  to the  $\phi(N_{pa})$  of a child  $N_{ch}$  of a sum node that led to  $S_j^{max}(D_j)$  and to  $\phi(N_{product})$  for all product nodes  $N_{product}$  along the path from the root to a leaf node
4 end
5 return  $\phi$ 
    
```

```

Algorithm 3: ComputeLabels
input :  $\tau = (G, \psi, \chi, NULL), \mathcal{D}, \phi, \lambda$ 
output : Labeled CSI-tree  $\tau = (G, \psi, \chi, \zeta)$ 
1 for Each edge  $e(N_{from}, N_{to})$  in  $G$  do
2   Compute class labels  $y$ 
3    $f = TrainModel(\mathcal{D}, \phi(N_{from}), y)$ 
4   Compute a set of important features  $I$  for  $f$ 
5    $\zeta(e) = \Lambda_{i \in I} i$  thresholded by  $\lambda$ 
6 end
7 return  $\tau = (G, \psi, \chi, \zeta)$ 
    
```

Empirical Evaluations



- ExSPN:
- recovers the CSIs encoded in an SPN – 83% of the CSIs recovered match the ground truth
 - extracts CSIs that are concise compared to the Apriori algorithm baseline – up to two orders of magnitude more concise
 - outputs CSI-tree on real-world clinical domain that is human-interpretable – validated by our medical expert Dr. Haas

Future Work

- Validating ExSPN on more relevant clinical studies
- Allowing domain experts to interact with the learned model
- Extending ExSPN to a broader class of models – TDPMs and beyond
- More types of explanations – beyond CSIs

Acknowledgement

The authors gratefully acknowledge the support of 1R01HD101246 from NICHD and Precision Health Initiative of Indiana University.

<https://starling.utdallas.edu/>
@STARLing_lab

