PADERBORN UNIVERSITY

# Neural Logic Reasoning

N'dah Jean Kouagou

DICE

Data Science Group
Paderborn University

September 03, 2021

- Motivation
- Model
  - Model Architecture – LINN
- Training
  - Loss Functions
  - Logical Regularizers and $L_2$ Regularization
- Evaluation
  - On Synthetic Data
  - On Recommendation Systems

# Motivation



- Generalization power of deep neural networks and reasoning capability of symbolic systems
- Solving real world tasks usually requires both reasoning and generalization capabilities

# Model Architecture (LINN)

- The authors proposed Logic Integrated Neural Network (LINN)
- The Model is dynamic

# Model Architecture (LINN)



- The authors proposed Logic Integrated Neural Network (LINN)
- The Model is dynamic

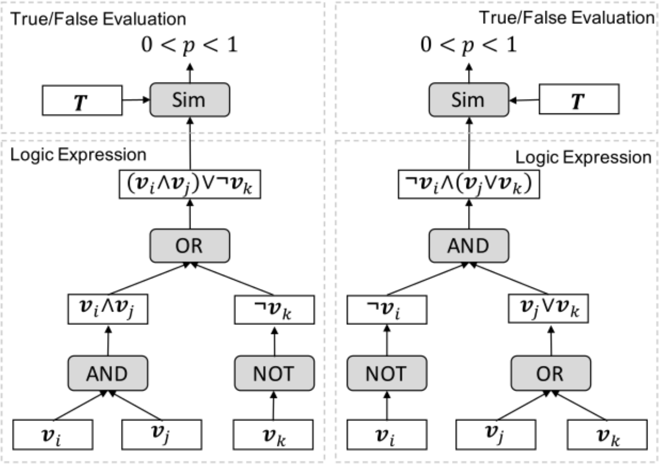


**(a)** $(v_i \wedge v_j) \vee \neg v_k$ **(b)** $\neg v_i \wedge (v_j \vee v_k)$

# Model Architecture (LINN)

- Logic variables are represented as vectors
- Each logic operation **AND**, **OR**, and **NOT** are represented as neural modules

# Model Architecture (LINN)



- Logic variables are represented as vectors
- Each logic operation **AND**, **OR**, and **NOT** are represented as neural modules
- The authors used MLP architecture for logic operations:

$$\textbf{AND}(w_i, w_j) = H_{a2} ReLU(H_{a1}(w_i \oplus w_j) + b_a)$$
$$\textbf{NOT}(w) = H_{n2} ReLU(H_{n1}(w) + b_n),$$

# Model Architecture (LINN)



- Logic variables are represented as vectors
- Each logic operation **AND**, **OR**, and **NOT** are represented as neural modules
- The authors used MLP architecture for logic operations:

$$\textbf{AND}(w_i, w_j) = H_{a2}ReLU(H_{a1}(w_i \oplus w_j) + b_a)$$
$$\textbf{NOT}(w) = H_{n2}ReLU(H_{n1}(w) + b_n),$$

where $H_{a1} \in \mathbb{R}^{d \times 2d}$, $H_{a2} \in \mathbb{R}^{d \times d}$, $b_a \in \mathbb{R}^d$, $H_{n1} \in \mathbb{R}^{d \times d}$, $H_{n2} \in \mathbb{R}^{d \times d}$, $b_n \in \mathbb{R}^d$, $\oplus$ is the concatenation operation, and $d$ is the embedding dimension.

# Loss Functions

- Two task-specific loss functions are considered:

$$L_t = L_{ce} = -\sum_i y_i \log(p_i) + (1 - y_i)\log(1 - pi) \quad (1)$$

$$L_t = L_{bpr} = -\sum_{e^+} \log(sigmoid(p(e^+) - p(e^-))) \quad (2)$$

- The first is the Cross-entropy loss, and the second is the Bayesian Personalized Ranking Loss

# Logic Regularizers and $L_2$ Regularization



| | Logical Rule | Equation | Logic Regularizer $r_i$ |
|---|---|---|---|
| NOT | Negation | $\neg T = F$ | $r_1 = \sum_{w \in W \cup \{T\}} Sim(\text{NOT}(\mathbf{w}), \mathbf{w})$ |
| | Double Negation | $\neg(\neg w) = w$ | $r_2 = \sum_{w \in W} 1 - Sim(\text{NOT}(\text{NOT}(\mathbf{w})), \mathbf{w})$ |
| AND | Identity | $w \wedge T = w$ | $r_3 = \sum_{w \in W} 1 - Sim(\text{AND}(\mathbf{w}, \mathbf{T}), \mathbf{w})$ |
| | Annihilator | $w \wedge F = F$ | $r_4 = \sum_{w \in W} 1 - Sim(\text{AND}(\mathbf{w}, \mathbf{F}), \mathbf{F})$ |
| | Idempotence | $w \wedge w = w$ | $r_5 = \sum_{w \in W} 1 - Sim(\text{AND}(\mathbf{w}, \mathbf{w}), \mathbf{w})$ |
| | Complementation | $w \wedge \neg w = F$ | $r_6 = \sum_{w \in W} 1 - Sim(\text{AND}(\mathbf{w}, \text{NOT}(\mathbf{w})), \mathbf{F})$ |
| OR | Identity | $w \vee F = w$ | $r_7 = \sum_{w \in W} 1 - Sim(\text{OR}(\mathbf{w}, \mathbf{F}), \mathbf{w})$ |
| | Annihilator | $w \vee T = T$ | $r_8 = \sum_{w \in W} 1 - Sim(\text{OR}(\mathbf{w}, \mathbf{T}), \mathbf{T})$ |
| | Idempotence | $w \vee w = w$ | $r_9 = \sum_{w \in W} 1 - Sim(\text{OR}(\mathbf{w}, \mathbf{w}), \mathbf{w})$ |
| | Complementation | $w \vee \neg w = T$ | $r_{10} = \sum_{w \in W} 1 - Sim(\text{OR}(\mathbf{w}, \text{NOT}(\mathbf{w})), \mathbf{T})$ |

# Logic Regularizers and $L_2$ Regularization



| | Logical Rule | Equation | Logic Regularizer $r_i$ |
|---|---|---|---|
| NOT | Negation | $\neg T = F$ | $r_1 = \sum_{w \in W \cup \{T\}} Sim(\text{NOT}(\mathbf{w}), \mathbf{w})$ |
| | Double Negation | $\neg(\neg w) = w$ | $r_2 = \sum_{w \in W} 1 - Sim(\text{NOT}(\text{NOT}(\mathbf{w})), \mathbf{w})$ |
| AND | Identity | $w \wedge T = w$ | $r_3 = \sum_{w \in W} 1 - Sim(\text{AND}(\mathbf{w}, \mathbf{T}), \mathbf{w})$ |
| | Annihilator | $w \wedge F = F$ | $r_4 = \sum_{w \in W} 1 - Sim(\text{AND}(\mathbf{w}, \mathbf{F}), \mathbf{F})$ |
| | Idempotence | $w \wedge w = w$ | $r_5 = \sum_{w \in W} 1 - Sim(\text{AND}(\mathbf{w}, \mathbf{w}), \mathbf{w})$ |
| | Complementation | $w \wedge \neg w = F$ | $r_6 = \sum_{w \in W} 1 - Sim(\text{AND}(\mathbf{w}, \text{NOT}(\mathbf{w})), \mathbf{F})$ |
| OR | Identity | $w \vee F = w$ | $r_7 = \sum_{w \in W} 1 - Sim(\text{OR}(\mathbf{w}, \mathbf{F}), \mathbf{w})$ |
| | Annihilator | $w \vee T = T$ | $r_8 = \sum_{w \in W} 1 - Sim(\text{OR}(\mathbf{w}, \mathbf{T}), \mathbf{T})$ |
| | Idempotence | $w \vee w = w$ | $r_9 = \sum_{w \in W} 1 - Sim(\text{OR}(\mathbf{w}, \mathbf{w}), \mathbf{w})$ |
| | Complementation | $w \vee \neg w = T$ | $r_{10} = \sum_{w \in W} 1 - Sim(\text{OR}(\mathbf{w}, \text{NOT}(\mathbf{w})), \mathbf{T})$ |

- The final loss for training LINN is:

$$L = L_t + \lambda_r \sum_i r_i + \lambda_w \sum_{w \in W} ||w||_F^2 + \lambda_\Theta ||\Theta||_F^2 \quad (3)$$

# Evaluation – On Synthetic Data

- Generate $n$ random logic variables and assign to each of them $T$ or $F$
- Create $m$ logic expressions from the generated variables

# Evaluation – On Synthetic Data

- Generate $n$ random logic variables and assign to each of them $T$ or $F$
- Create $m$ logic expressions from the generated variables

$$(\neg v_{80} \wedge v_{56} \wedge v_{71}) \vee (\neg v_{46} \wedge \neg v_7 \wedge v_{51} \wedge \neg v_{47} \wedge v_{26}) \vee v_{45} \vee (v_{31} \wedge v_{15} \wedge v_2 \wedge v_{46}) = T$$

$$(\neg v_{19} \wedge \neg v_{65}) \vee (v_{65} \wedge \neg v_{24} \wedge v_9 \wedge \neg v_{83}) \vee (\neg v_{48} \wedge \neg v_9 \wedge \neg v_{51} \wedge v_{75}) = F$$

$$\neg v_{98} \vee (\neg v_{76} \wedge v_{66} \wedge v_{13}) \vee (v_{97} \wedge v_{89} \wedge v_{45} \wedge v_{83}) = T$$

$$(v_{43} \wedge v_{21} \wedge \neg v_{53}) = F$$

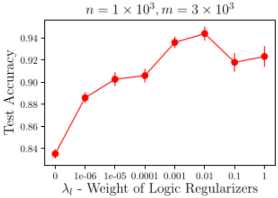


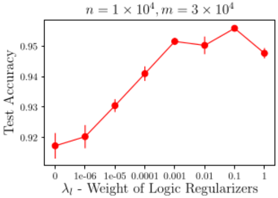

# Evaluation – On Synthetic Data



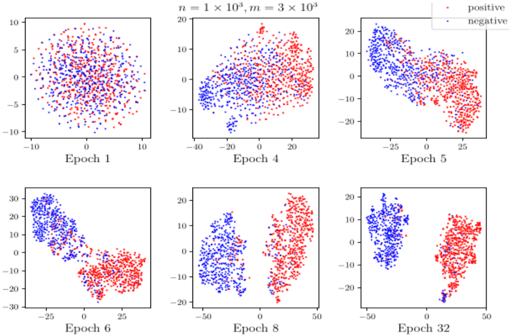


| | $n = 1 \times 10^3, m = 3 \times 10^3$ | | $n = 1 \times 10^4, m = 3 \times 10^4$ | |
|---|---|---|---|---|
| | Accuracy | RMSE | Accuracy | RMSE |
| Bi-RNN [32] | 0.6493 ± 0.0102 | 0.4736 ± 0.0032 | 0.6942 ± 0.0028 | 0.4492 ± 0.0009 |
| Bi-LSTM [11] | 0.5933 ± 0.0107 | 0.5181 ± 0.0162 | 0.6847 ± 0.0051 | 0.4494 ± 0.0020 |
| CNN [19] | 0.6380 ± 0.0043 | 0.5085 ± 0.0158 | 0.6787 ± 0.0025 | 0.4557 ± 0.0016 |
| LINN-$R_l$ | *0.8353 ± 0.0043* | *0.3880 ± 0.0069* | *0.9173 ± 0.0042* | *0.2733 ± 0.0065* |
| LINN | **0.9440 ± 0.0064*** | **0.2318 ± 0.0124*** | **0.9559 ± 0.0006*** | **0.2081 ± 0.0018*** |

* Significantly better than the best of the other results (italic ones) with $p < 0.05$

# Recommendation Systems



| | ML-100k | | | Amazon Electronics | | |
|---|---|---|---|---|---|---|
| | nDCG@10 | Hit@1 | time/epoch | nDCG@10 | Hit@1 | time/epoch |
| BPRMF [31] | 0.3664 ± 0.0017 | 0.1537 ± 0.0036 | 4.9s | 0.3514 ± 0.0002 | 0.1951 ± 0.0004 | 112.1s |
| SVD++ [21] | 0.3675 ± 0.0024 | 0.1556 ± 0.0044 | 30.4s | 0.3582 ± 0.0004 | 0.1930 ± 0.0006 | 469.3s |
| STAMP [25] | 0.3943 ± 0.0016 | 0.1706 ± 0.0022 | 8.3s | 0.3954 ± 0.0003 | 0.2215 ± 0.0003 | 352.7s |
| GRU4Rec [16] | 0.3973 ± 0.0016 | 0.1745 ± 0.0038 | 7.1s | 0.4029 ± 0.0009 | 0.2262 ± 0.0009 | 225.0s |
| NARM [24] | *0.4022 ± 0.0015* | 0.1771 ± 0.0016 | 9.6s | 0.4051 ± 0.0006 | 0.2292 ± 0.0005 | 268.8s |
| LINN-$R_l$ | 0.4022 ± 0.0027 | *0.1783 ± 0.0043* | 20.7s | *0.4152 ± 0.0014* | *0.2396 ± 0.0019* | 498.0s |
| LINN | **0.4064 ± 0.0015*** | **0.1850 ± 0.0053*** | 30.7s | **0.4191 ± 0.0012*** | **0.2438 ± 0.0014*** | 754.9s |

* Significantly better than the best of other results (italic ones) with $p < 0.05$

9 random movies ranked on popularity

$Sim(\mathbf{v}_i \wedge \mathbf{v}_j, \mathbf{T})$