36th AAAI Conference on Artificial Intelligence (AAAI-22),
February 22-March 1, 2022,
VIRTUAL CONFERENCE

Fully Adaptive Framework: Neural Computerized Adaptive Testing for Online Education

Yan Zhuang[1], Qi Liu[1*], Zhenya Huang[1], Zhi Li[1], Shuanghong Shen[1], Haiping Ma[2]
[1]Anhui Province Key Laboratory of Big Data Analysis and Application, School of Data Science & School of Computer Science and Technology, University of Science and Technology of China (USTC); [2]Anhui University;

Reporter: Yan Zhuang

AAAI-2022

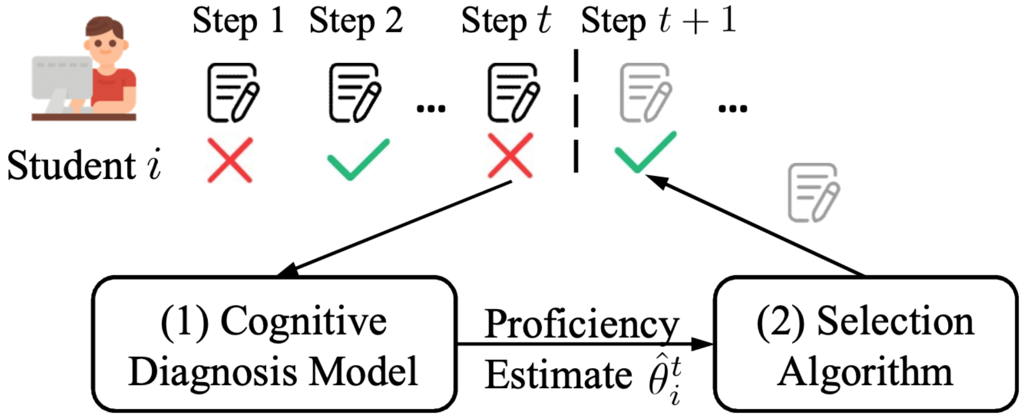

中国科学技术大学
University of Science and Technology of China

Background

Typical CAT procedure.

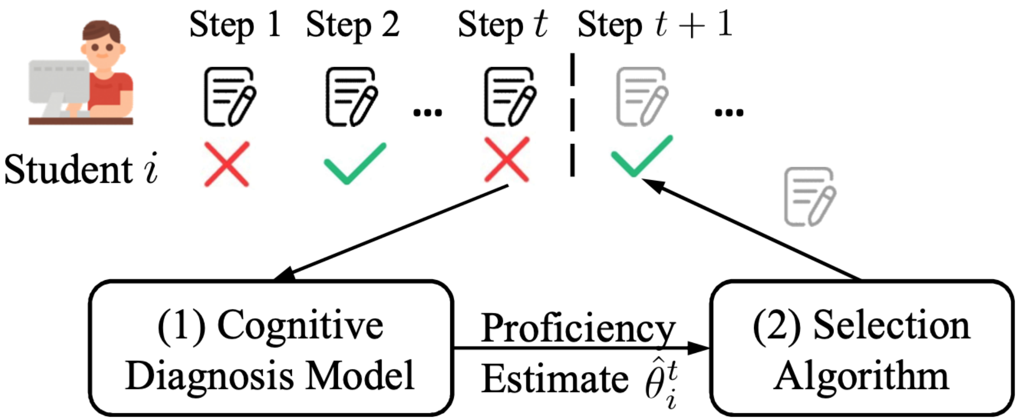

Goal:

- Measuring student's proficiency accurately
- Reducing test length

- **(1) Cognitive Diagnosis Model (CDM),** which first estimates the student's current proficiency $\hat{\theta}^t$ based on previous t responses. The representative models are IRT and NCDM.

- **(2) The Selection Algorithm** then selects the next question, guided by his/her current proficiency estimate $\hat{\theta}^t$ above. Most algorithms are **model-specific**, which are specially designed by experts according to different CDMs' characteristics.

Outline

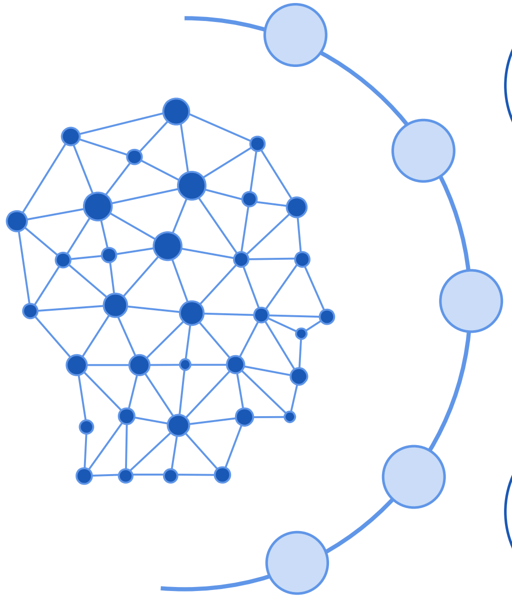

Outline

1 Background

2 Formulation

3 Implementation

4 Experiment

5 Conclusion

Background

How to **accurately** and **efficiently** measure student's ability/proficiency?

- Paper-and-pencil Examination
  - Too many questions - inefficient and boring
  - Fixed time/place - inflexible

- Computerized Adaptive Testing (CAT)
  - Personalization and reduce test length
  - Flexible time/place

Background

Typical CAT procedure.

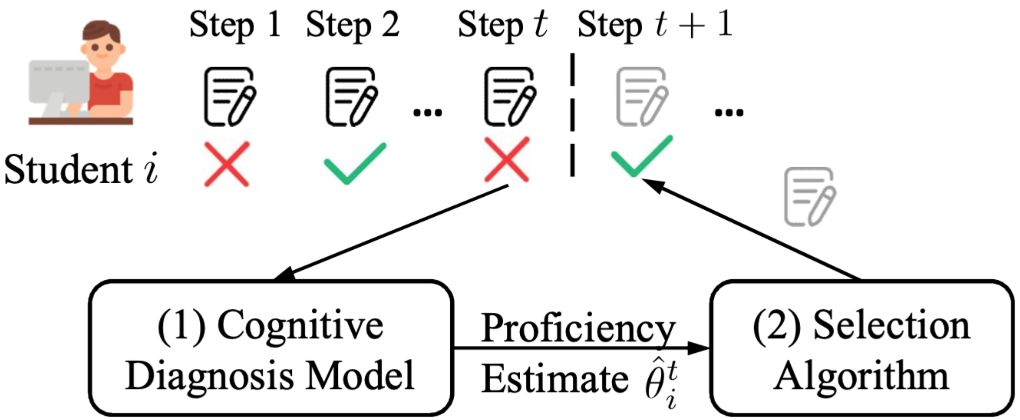

Goal:

- Measuring student's proficiency accurately
- Reducing test length

- **(1) Cognitive Diagnosis Model (CDM),** which first estimates the student's current proficiency $\hat{\theta}^t$ based on previous t responses. The representative models are IRT and NCDM.

- **(2) The Selection Algorithm** then selects the next question, guided by his/her current proficiency estimate $\hat{\theta}^t$ above. Most algorithms are **model-specific**, which are specially designed by experts according to different CDMs' characteristics.

Background

Typical CAT procedure.

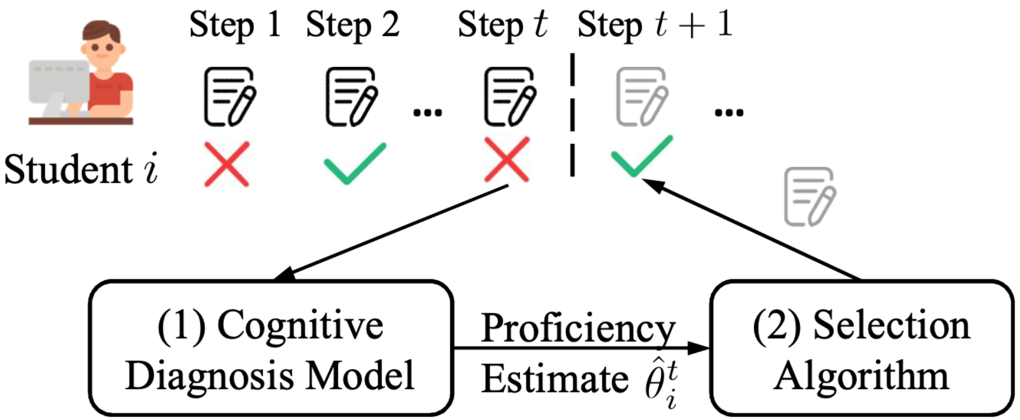

Goal:

- Measuring student's proficiency accurately
- Reducing test length

Limited adaptability:

- For **student**, selection algorithm's efficiency heavily relies on the accuracy of $\hat{\theta}^t$. → Poor robustness, Loss of information
- For **CDM**, have to understand how a specific CDM works to design the matched algorithms. → Model-specific, Labor-intensive
- For **questions**, such pre-defined algorithms have individual "preference" in selections (e.g., MFI). → Exposure unbalance, Test insecur

Outline

Outline

1 Background

2 Formulation

3 Implementation

4 Experiment

5 Conclusion

Formulation

- **The Learnable Selection Algorithm π:** → Using large-scale response data
  - **Bi-level Optimization**

The responses of student i is randomly divided into:

(1) Support set $\mathcal{D}_s^i$ (2) Query set $\mathcal{D}_u^i$

Outer-level

$$\pi^* = \arg\min_{\pi} \frac{1}{n} \sum_{i=1}^{n} \sum_{t=1}^{T} \frac{1}{|\mathcal{D}_u^i|} \sum_{(q,a)\in\mathcal{D}_u^i} l(a, \mathcal{M}(q|\hat{\theta}_i^t)),$$

Inner-level

$$\text{s.t.} \quad \hat{\theta}_i^t = \arg\min_{\theta_i} \sum_{(q,a)\in\mathcal{D}_s^{i(t)}} l\left(a, \mathcal{M}\left(q|\theta_i\right)\right),$$

$$\text{where} \quad \mathcal{D}_s^{i(t)} = \{q_1, a_{i(1)}, \ldots, q_t, a_{i(t)}\} \quad \text{and}$$

$$q_t \sim \pi\left(q_1, a_{i(1)}, \ldots, q_{t-1}, a_{i(t-1)}\right).$$

Formulation

- **The Learnable Selection Algorithm π:**

Using large-scale response data

- **Bi-level Optimization**

The responses of student *i* is randomly divided into:

(1) Support set $\mathcal{D}_s^i$ **(2) Query set $\mathcal{D}_u^i$**

Outer-level

$$\pi^* = \arg\min_{\pi} \frac{1}{n} \sum_{i=1}^{n} \sum_{t=1}^{T} \frac{1}{|\mathcal{D}_u^i|} \sum_{(q,a)\in\mathcal{D}_u^i} l(a, \mathcal{M}(q|\hat{\theta}_i^t)),$$

Proficiency estimate on support set

Inner-level

$$\text{s.t. } \boxed{\hat{\theta}_i^t} = \arg\min_{\theta_i} \sum_{(q,a)\in\mathcal{D}_s^{i(t)}} l\left(a, \mathcal{M}\left(q|\theta_i\right)\right),$$

$$\text{where } \quad \mathcal{D}_s^{i(t)} = \{q_1, a_{i(1)}, ..., q_t, a_{i(t)}\} \quad \text{and}$$

$$q_t \sim \pi\left(q_1, a_{i(1)}, ..., q_{t-1}, a_{i(t-1)}\right).$$

Formulation

- **The Learnable Selection Algorithm π:** → Using large-scale response data
  - **Bi-level Optimization**

The responses of student *i* is randomly divided into:

(1) Support set $\mathcal{D}_s^i$ (2) Query set $\mathcal{D}_u^i$

Sum all the test steps and students

Outer-level

$$\pi^* = \arg\min_{\pi} \frac{1}{n} \boxed{\sum_{i=1}^{n} \sum_{t=1}^{T}} \frac{1}{|\mathcal{D}_u^i|} \sum_{(q,a)\in\mathcal{D}_u^i} l(a, \mathcal{M}(q|\boxed{\hat{\theta}_i^t})),$$

Fit of estimate on query set

Proficiency estimate on support set

Inner-level

$$\text{s.t.} \ \boxed{\hat{\theta}_i^t} = \arg\min_{\theta_i} \sum_{(q,a)\in\mathcal{D}_s^{i(t)}} l\left(a, \mathcal{M}\left(q|\theta_i\right)\right),$$

$$\text{where} \quad \mathcal{D}_s^{i(t)} = \{q_1, a_{i(1)}, \ldots, q_t, a_{i(t)}\} \quad \text{and}$$

$$q_t \sim \pi\left(q_1, a_{i(1)}, \ldots, q_{t-1}, a_{i(t-1)}\right).$$

Formulation

- **The Learnable Selection Algorithm π:**
  - **Reinforcement Learning Formulation**

$$
\begin{aligned}
& \min_{\pi} \frac{1}{n} \sum_{i=1}^{n} \sum_{t=1}^{T} \frac{1}{|\mathcal{D}_u^i|} \sum_{(q,a)\in\mathcal{D}_u^i} l(a, \mathcal{M}(q|\hat{\theta}_i^t)) \\
\triangleq \quad & \max_{\pi} \mathbb{E}_{i\sim\pi} \left[\sum_{t=1}^{T} -\frac{1}{|\mathcal{D}_u^i|} \sum_{(q,a)\in\mathcal{D}_u^i} l(a, \mathcal{M}(q|\hat{\theta}_i^t)) \right] \\
= \quad & \max_{\pi} \mathbb{E}_{i\sim\pi} \left[\sum_{t=1}^{T} -\mathcal{L}_{\mathcal{M}}(\mathcal{D}_u^i, \hat{\theta}_i^t) \right],
\end{aligned}
$$

Formulation

◆ The Learnable Selection Algorithm π:

◆ Reinforcement Learning Formulation

$$\min_{\pi} \frac{1}{n} \sum_{i=1}^{n} \sum_{t=1}^{T} \frac{1}{|\mathcal{D}_u^i|} \sum_{(q,a)\in\mathcal{D}_u^i} l(a, \mathcal{M}(q|\hat{\theta}_i^t))$$

$$\triangleq \max_{\pi} \mathbb{E}_{i\sim\pi} \left[\sum_{t=1}^{T} -\frac{1}{|\mathcal{D}_u^i|} \sum_{(q,a)\in\mathcal{D}_u^i} l(a, \mathcal{M}(q|\hat{\theta}_i^t)) \right]$$

$$= \max_{\pi} \mathbb{E}_{i\sim\pi} \left[\sum_{t=1}^{T} \underbrace{-\mathcal{L}_{\mathcal{M}}(\mathcal{D}_u^i, \hat{\theta}_i^t)}_{\text{Reward}} \right],$$

- **State:** previous t responses

$$s_t = \{q_1, a_{i(1)}, ..., q_{t-1}, a_{i(t-1)}\}.$$

- **Action:** selection of next question: q_t

- **Transition:** the uncertainty comes from next question's correction $a_{i(t)}$

$$P(s_{t+1}|s_t, q_t)$$

- **Reward:** negative loss of the estimated proficiency of on query set

$$-\mathcal{L}_{\mathcal{M}}(\mathcal{D}_u^i, \hat{\theta}_i^t)$$

Formulation

◆ The Learnable Selection Algorithm π:

◆ Reinforcement Learning Formulation

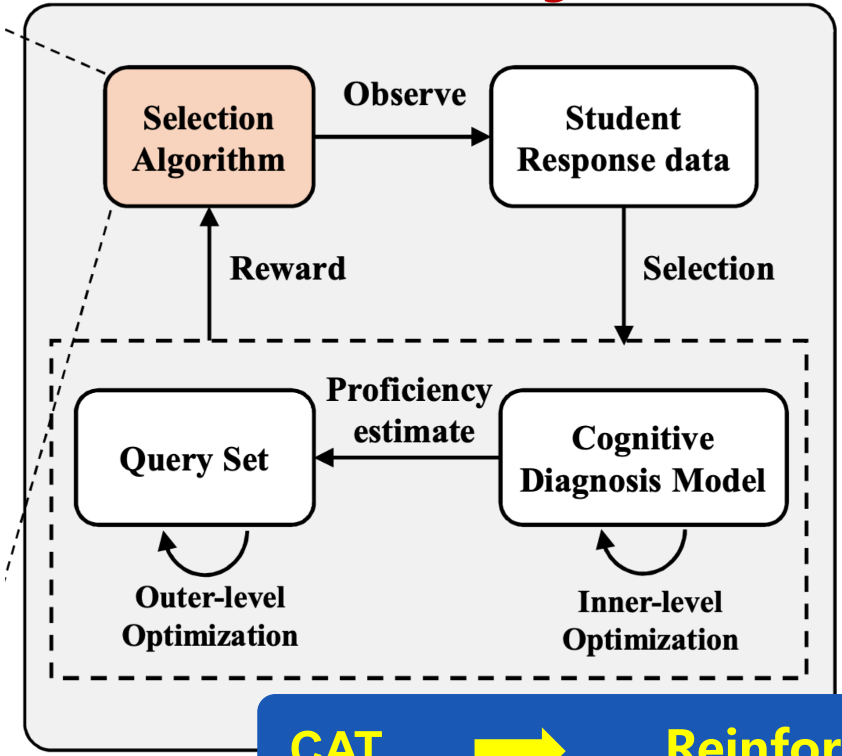

- **State:** previous t responses

$$s_t = \{q_1, a_{i(1)}, \ldots, q_{t-1}, a_{i(t-1)}\}.$$

- **Action:** selection of next question: q_t
- **Transition:** the uncertainty comes from next question's correction $a_{i(t)}$

$$P(s_{t+1}|s_t, q_t)$$

- **Reward:** negative loss of the estimated proficiency of on query set

$$-\mathcal{L}_{\mathcal{M}}(\mathcal{D}_u^i, \hat{\theta}_i^t)$$

CAT ➡ Reinforcement Learning problem

Outline

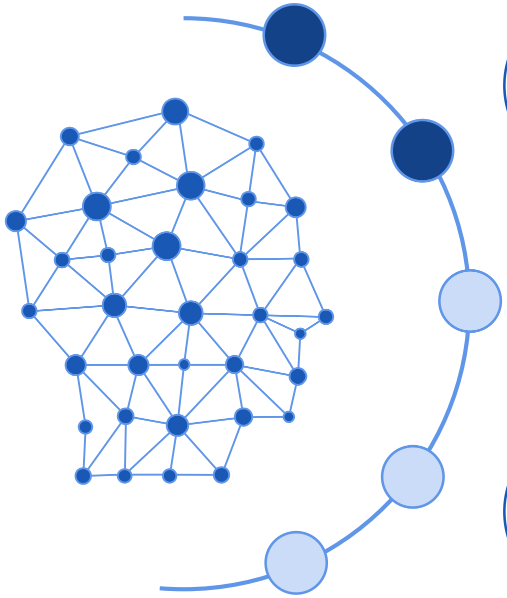

1. Background
2. Formulation
3. Implementation
4. Experiment
5. Conclusion

Outline

Methodology

◆ Selection algorithm in our NCAT framework:

- ◆ Double-Channel Performance Learning
- ◆ Contradiction Learning

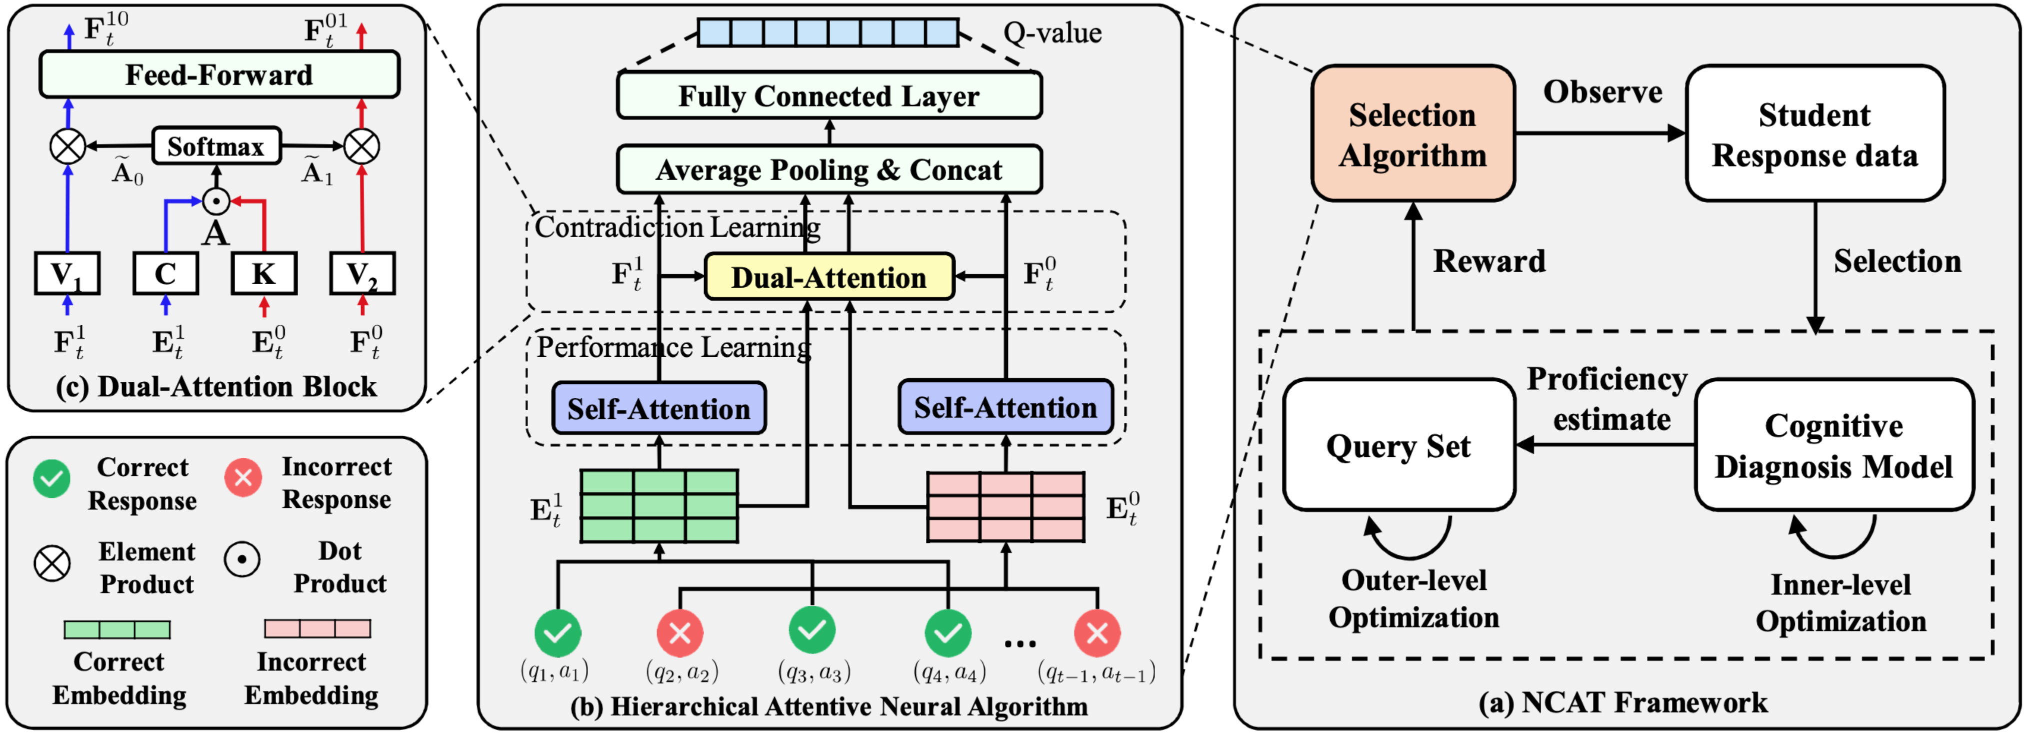

Methodology

◆ Selection algorithm in our NCAT framework:

◆ Double-Channel Performance Learning

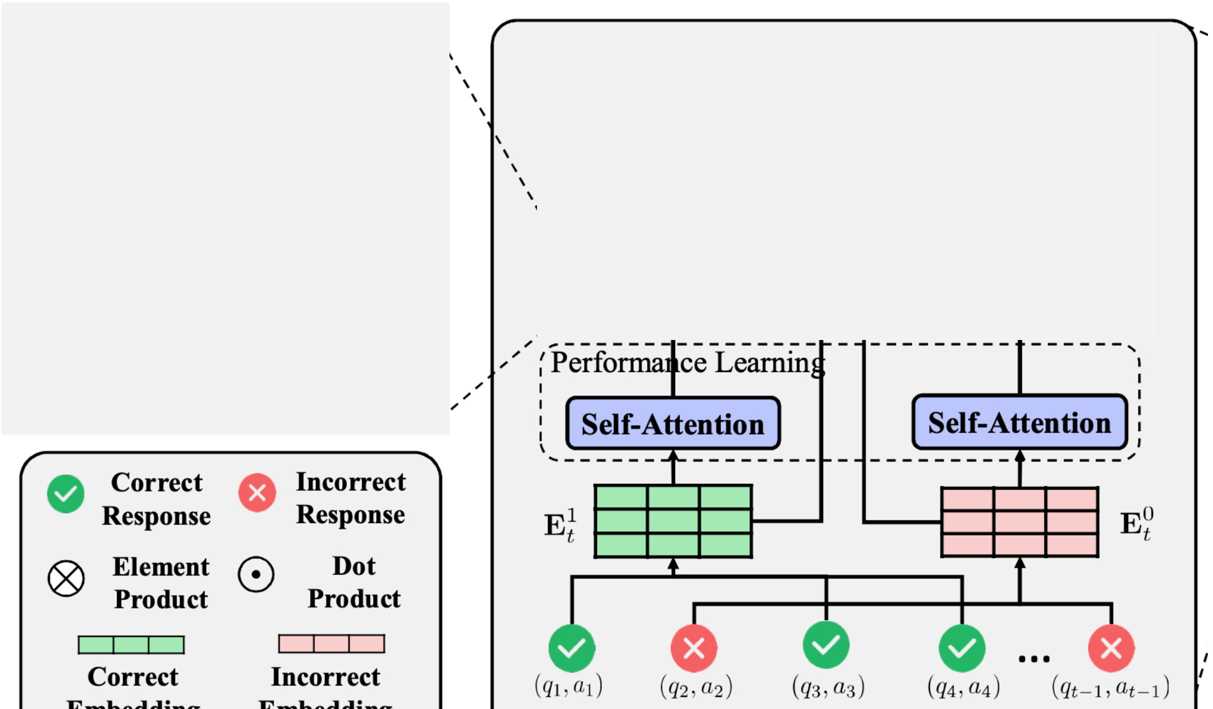

(b) Hierarchical Attentive Neural Algorithm

Correct and incorrect responses are imbalanced:

incorrect < correct

Separately capture the information of versatile student performance (double-channel)

Methodology

◆ Selection algorithm in our NCAT framework:

◆ Contradiction Learning

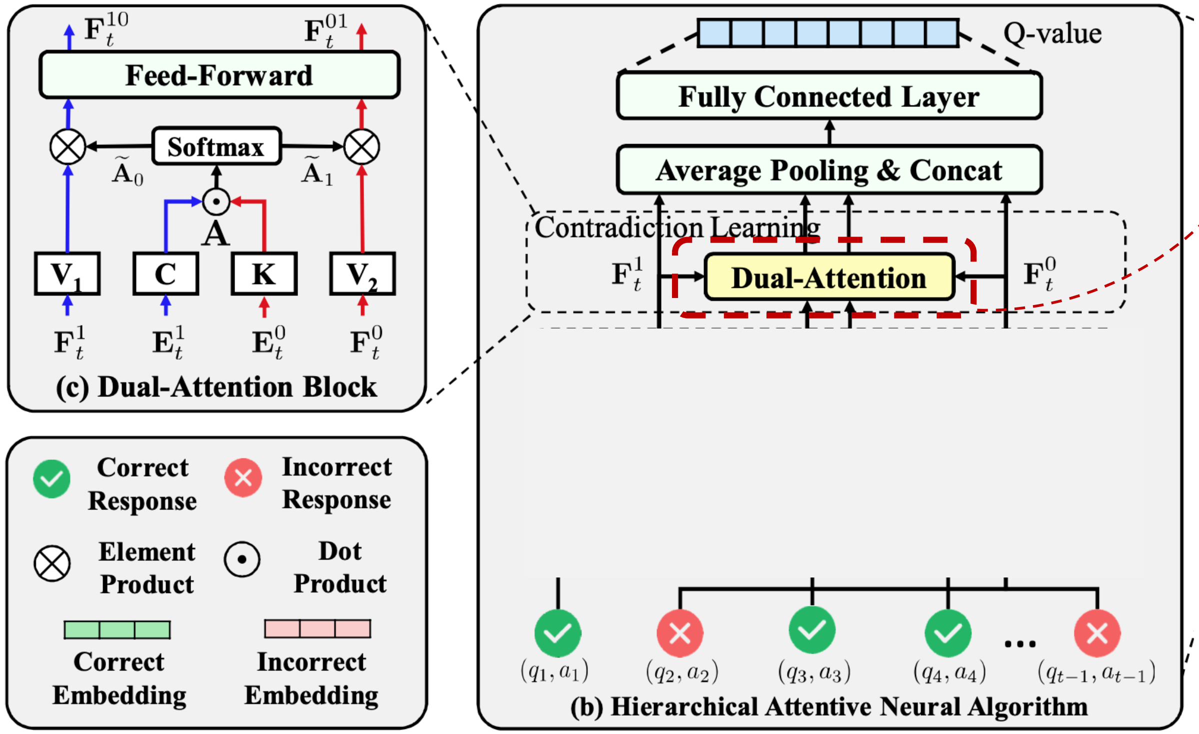

Guess and Slip factors

Contradictions between the correct and incorrect responses:

harder simpler

Possible Contradiction:

- **Guess** factor in the harder
- **Slip** factor in the simpler
- **Both**

Outline

Outline

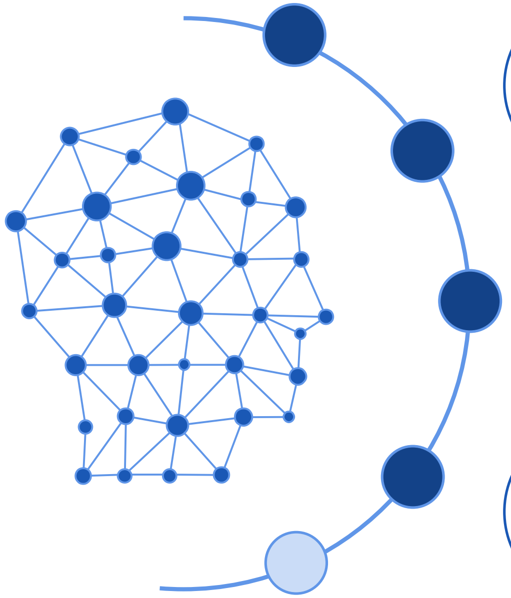

1. Background
2. Formulation
3. Implementation
4. Experiment
5. Conclusion

Experiment

Setups

- **Dataset**
  - Real-world datasets from three **online tutoring system**
  - Involving two classic CDM: **IRT and NCDM**
- **Comparison Methods**
  - **Traditional information/uncertainty-based:**
    - MFI, KLI
    - MAAT (active learning)
  - **Meta Learning :**
    - BOBCAT
- **Evaluation Metrics:** AUC, ACC, MSE

Experiment

Results

Our proposed NCAT achieves the best performance on all datasets and all types of CDMs.

Dataset	ASSIST						NIPS-EDU						EXAM					
CDM	IRT			NCDM			IRT			NCDM			IRT			NCDM		
Metric	ACC						ACC						ACC					
Step	5	10	20	5	10	20	5	10	20	5	10	20	5	10	20	5	10	20
RAND	0.7099	0.7168	0.7241	0.7036	0.7096	0.7205	0.6290	0.6583	0.6868	0.6212	0.6632	0.6921	0.7212	0.7540	0.8100	0.7230	0.7618	0.8144
MFI	0.7224	0.7296	0.7412	–	–	–	0.6464	0.6765	0.7056	–	–	–	0.7495	0.7753	0.8382	–	–	–
KLI	0.7230	0.7298	0.7418	–	–	–	0.6456	0.6719	0.7016	–	–	–	0.7556	0.7841	0.8431	–	–	–
MAAT	0.7233	0.7291	0.7422	0.7268	0.7314	0.7487	0.6481	0.6742	0.7134	0.6432	0.6811	0.7180	0.7590	0.7972	0.8450	0.7623	0.8003	0.8455
BOBCAT	0.7262	0.7328	0.7488	0.7298	0.7409	0.7492	0.6559	0.6812	0.7225	0.6630	0.6952	0.7236	0.7663	0.7991	0.8444	0.7712	0.8100	0.8442
NCAT	**0.7330**	**0.7478**	**0.7562**	**0.7354**	**0.7539**	**0.7564**	**0.6606**	**0.7032**	**0.7321**	**0.6742**	**0.7159**	**0.7334**	**0.7813**	**0.8165**	**0.8516**	**0.7837**	**0.8235**	**0.8546**
Metric	AUC						AUC						AUC					
Step	5	10	20	5	10	20	5	10	20	5	10	20	5	10	20	5	10	20
RAND	0.6884	0.6978	0.7102	0.6891	0.6971	0.7191	0.6550	0.6872	0.7229	0.6591	0.6988	0.7260	0.6714	0.6923	0.7683	0.6791	0.7021	0.7795
MFI	0.6992	0.7100	0.7297	–	–	–	0.6728	0.7066	0.7458	–	–	–	0.6950	0.7188	0.7834	–	–	–
KLI	0.7004	0.7104	0.7316	–	–	–	0.6706	0.7026	0.7390	–	–	–	0.7056	0.7293	0.7869	–	–	–
MAAT	0.7112	0.7122	0.7350	0.7135	0.7211	0.7444	0.6729	0.7034	0.7481	0.6713	0.7139	0.7473	0.7013	0.7345	0.7910	0.7023	0.7376	0.7984
BOBCAT	0.7155	0.7200	0.7423	0.7181	0.7356	0.7463	0.6841	0.7100	0.7571	0.6910	0.7204	0.7622	0.7067	0.7374	0.7908	0.7107	0.7421	0.7921
NCAT	**0.7187**	**0.7318**	**0.7546**	**0.7208**	**0.7391**	**0.7525**	**0.6889**	**0.7307**	**0.7600**	**0.7043**	**0.7394**	**0.7661**	**0.7124**	**0.7489**	**0.8022**	**0.7131**	**0.7589**	**0.8152**

Experiment

Contradiction in Responses

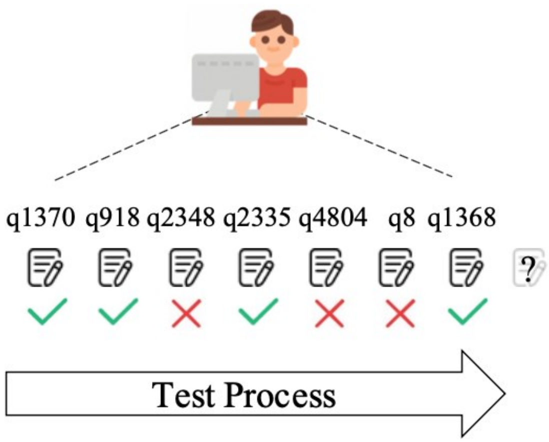

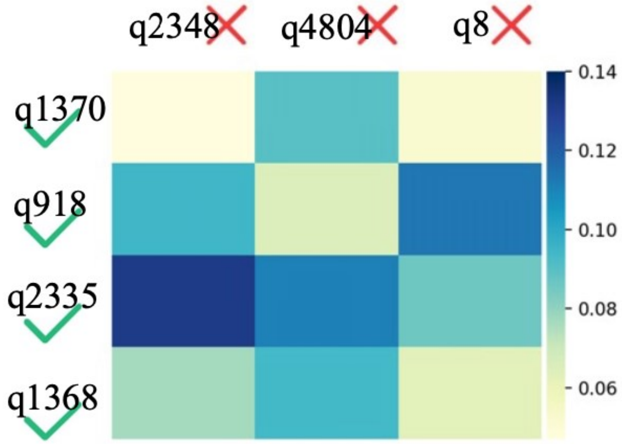

Correct	Related Concepts	Difficulty
q1370	Unary Linear Inequality	0.069
q918	Real Number	0.150
q2335	Linear Equations in Three Variables	0.171
q1368	Cube Root	0.140

Incorrect	Related Concepts	Difficulty
q2348	Linear Equation in One Variable	0.151
q4804	Function	0.164
q8	Real Number	0.154

Experiment

Contradiction in Responses

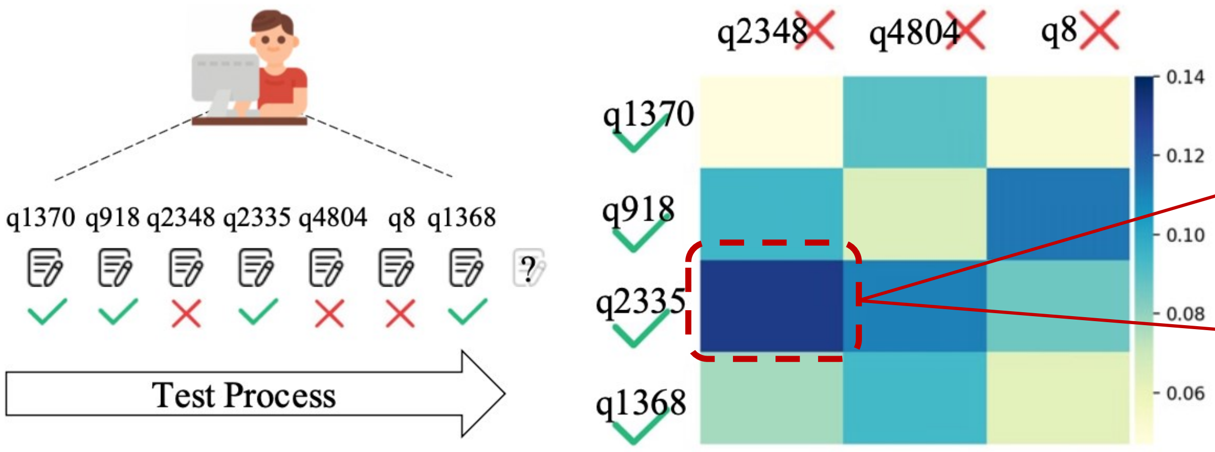

Correct	Related Concepts	Difficulty
q1370	Unary Linear Inequality	0.069
q918	Real Number	0.150
q2335	Linear Equations in Three Variables	0.171
q1368	Cube Root	0.140

Incorrect	Related Concepts	Difficulty
q2348	Linear Equation in One Variable	0.151
q4804	Function	0.164
q8	Real Number	0.154

Experiment

Contradiction in Responses

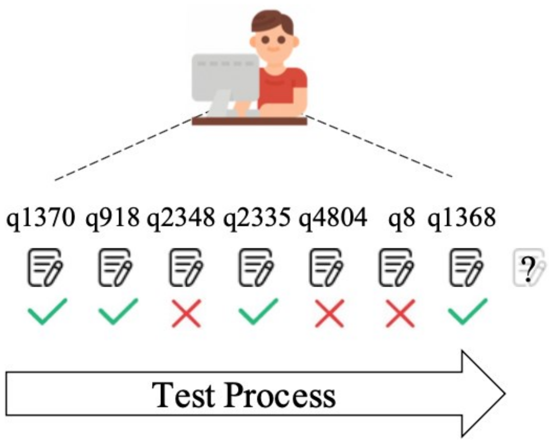

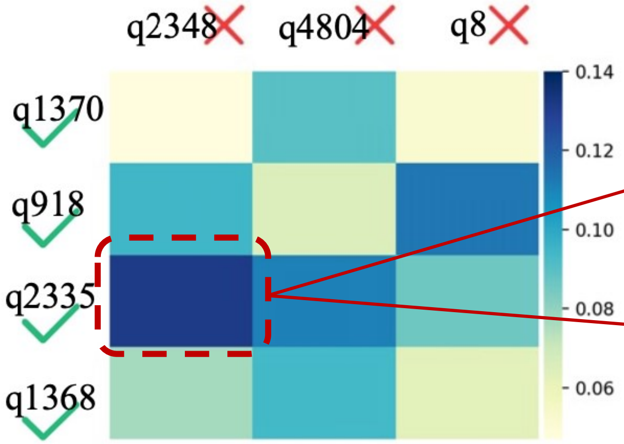

Correct	Related Concepts	Difficulty
q1370	Unary Linear Inequality	0.069
q918	Real Number	0.150
q2335	Linear Equations in Three Variables	0.171
q1368	Cube Root	0.140

Incorrect	Related Concepts	Difficulty
q2348	Linear Equation in One Variable	0.151
q4804	Function	0.164
q8	Real Number	0.154

He answers a difficult question (q2335) correctly, but get a simple one wrong (q2348)

These observations imply that NCAT provides a good way to capture the complex relationship in questions-students for better selections.

Outline

Outline

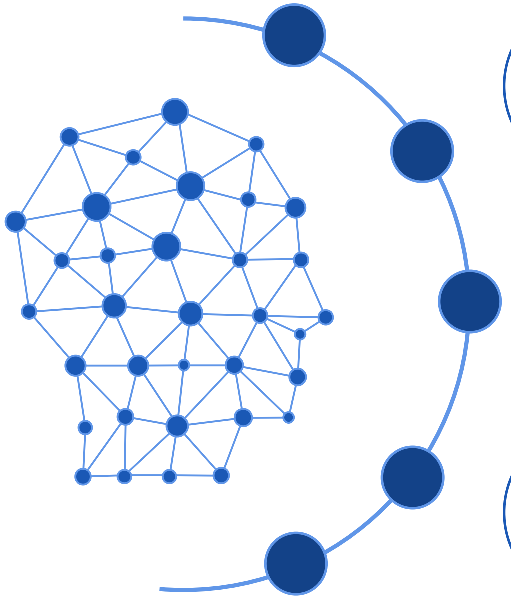

1. Background
2. Formulation
3. Implementation
4. Experiment
5. Conclusion

Conclusion

Conclusion

- **Formally redefine CAT as the Reinforcement Learning problem.**
  - No need to design the selection algorithm manually
- **Propose a novel Contradiction Learning module to model complex interactions**
  - Capture the Guess and Slip factors
- **Conduct extensive experiments with real-world educational datasets**
  - Efficiency, Robustness, Diversity, Exposure Rate

Thanks for your listening!

For more details, please refer to our paper!

Reporter: Yan Zhuang
zykb@mail.ustc.edu.cn

中国科学技术大学
University of Science and Technology of China