

Eedi and the NeurIPS 2020 Education Challenge Dataset

Simon Woodhead

I= Eedi

Why should I use the NeurIPS 2020 Education Challenge Dataset?

- Free, large-scale, high-quality edu dataset
- Taken from use across the world and over multiple years
- Proven fruitful for various tasks and models
- Fully documented in a white paper
- Evaluation scripts and numerous benchmarks
- Eedi is keen and ready to support

Novel bits:

- Crowd-sourced questions from teachers
- Capturing student misconceptions



Eedi who?

Where does this data come from?


Eedi is an edtech startup from the United Kingdom

- Online education platform co-founded by a data scientist and a (pretty famous) maths teacher
- Goal to improve learning through a better understanding of students' misconceptions
- Crowd-sourcing from teachers
- Track record of connecting research to product for real world impact
- Pedagogical weapon of choice are diagnostic questions...



What is a Diagnostic Question?

A multiple-choice question where distractors are key



What is the size of the angle marked P?

- A 150°
- B 65°
- C 115°
- D 85°



The Dataset

Show me the
~~money~~ number of
data points

- 125k Students
- 28k Questions
- 20M Answers (A, B, C or D)



Who are these students?

- Age
- Gender
- Eligibility for free school meals
- All data is completely anonymised

What are these questions about?

- Subjects (e.g. “algebra”)
- Topics (e.g. “factorizing”)
- Subtopics (e.g. “factorizing into a single bracket”)
- Topic hierarchy describing the taxonomy

Tell me more about the answers

- When the student answered the question
- Student's confidence in their answer
- Which quiz the question was in
- Which class the quiz was assigned to
- Whether the assignment was part of a longer scheme of work



Microsoft Research and Product Integration

What have you done so far?

- Worked with Microsoft Research to develop a model for imputation and active learning
- Active learning used in Eedi to choose the most informative question for a student
- Reduces number of questions needed to “level” a student, replaces manual onboarding
- Imputation is then used to calculate a knowledge “score” which can be tracked over time

Do you use Reinforcement Learning?

- No, the model, called EDDI (Ma et al, 2018), is Bayesian experimental design based
- Compared to Deep Reinforcement Learning (Zannone et al., 2019)
- RL is nonmyopic and it can perform better compared to EDDI but it is much less sample efficient
- We welcome any approaches to personalized decision making!

The NeurIPS 2020 Education Challenge

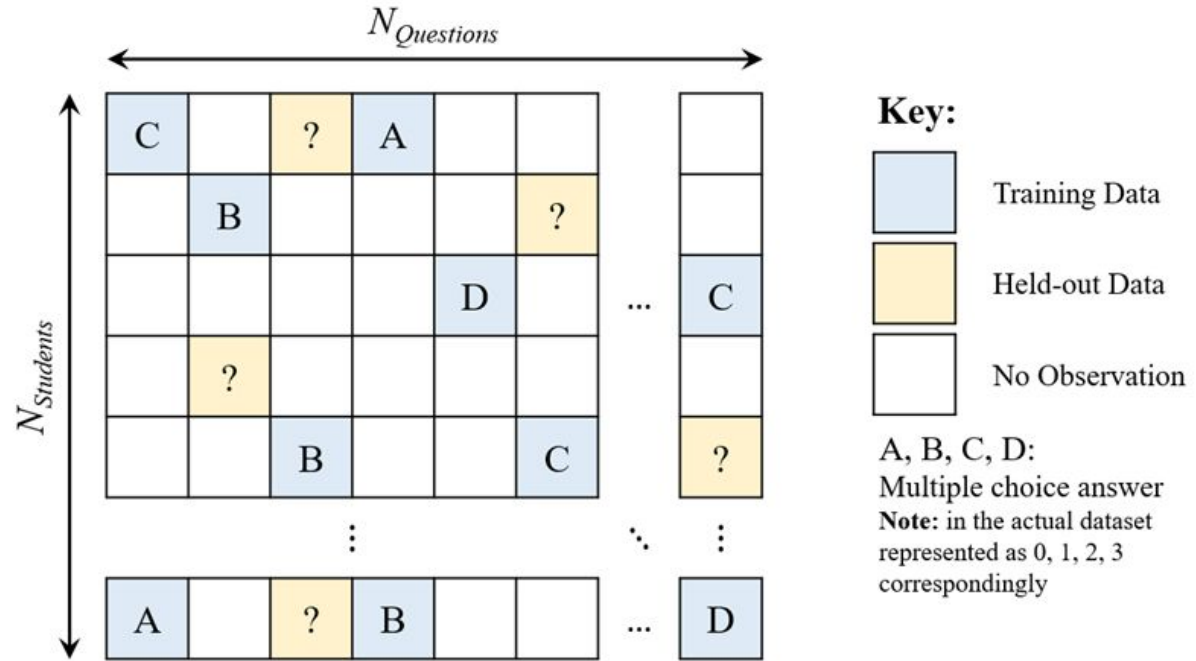
What was this
competition
about?

Four tasks designed to cover a range of real-world
education challenges:

1. Predicting student responses (binary)
2. Predicting student responses (ABCD)
3. Question quality assessment
4. Personalized questions



What is the real-world impact of Task 2?



- Predicting multiple-choice responses suggests specific misconceptions

Who participated?

- Most popular competition at NeurIPS 2020
- 382 teams
- 3696 entries
- Extended abstracts and code released on the competition website

**A Meta-learning Framework for Personalized
Question Selection**

<https://eedi.com/projects/neurips-education-challenge>



Did any participants use reinforcement learning?

Yes! On Task 4:
Personalized Questions

- Dynamically generate a personalized sequence of questions to maximise our understanding of a particular student
- The winning team tried actor-critic RL before settling on a fixed active learning strategy

Nevertheless, learning a robust model from limited observation for each student (irrespective of the number of students) remains a key challenge in this framework. We observed that the actor-critic algorithm [9] performs similarly to active learning-based methods in our preliminary experiments with a larger test set (20% students). We believe that learning a policy (using reinforcement learning) might improve results than a fixed active learning algorithm with more training and enough hyperparameter tuning.

[A Meta-learning Framework for Personalized Question Selection](#)

(Ghosh and Lan)



Conclusion

OK. I'm sold. How do I get started?

Download the dataset!

- There is a dedicated [dataset webpage](#)
- A [white paper](#) describing the dataset and the NeurIPS competition tasks
- Get in touch with me
simon.woodhead@eedi.co.uk

<https://eedi.com/projects/neurips-education-challenge>



Thanks

Eedi Team

Microsoft Research
Team

NeurIPS Competition
Organisers

NeurIPS Competition Zichao Wang, Angus Lamb, Evgeny Saveliev, Pashmina Cameron, Yordan Zaykov, José Miguel Hernández-Lobato, Richard E. Turner, Richard G. Baraniuk, Craig Barton, Simon Peyton Jones, Cheng Zhang

Microsoft Research Miltos Allamanis, Javier González Hernández, Sebastian Nowozin, Jooyeon Kim, Pablo Morales Alvarez, Virginia Aglietti

Eedi Digory Smith, Ben Caulfield

