HUB Român de Inteligență Artificială

UNIVERSITATEA TEHNICĂ DIN CLUJ-NAPOCA

Liquid Neural Networks for Adaptive Difficulty Selection in Personalized E-Learning

25th RoEduNet Conference: Networking in Education and Research
Iași, 19–21 November 2026

Sonia Grigor, Constantin Nandra, Dorian Gorgan

Computer Science Department
Technical University of Cluj-Napoca, Romania

Agenda

1. The problem and the AdaptEd idea
2. The LTC classifier and its data
3. Results: LTC vs. rules
4. Simulation, takeaways, next steps

The problem

2σ

One-to-one tutoring beats the classroom by 2 SD (Bloom)

But tutoring doesn't scale — and adaptive systems still fall short

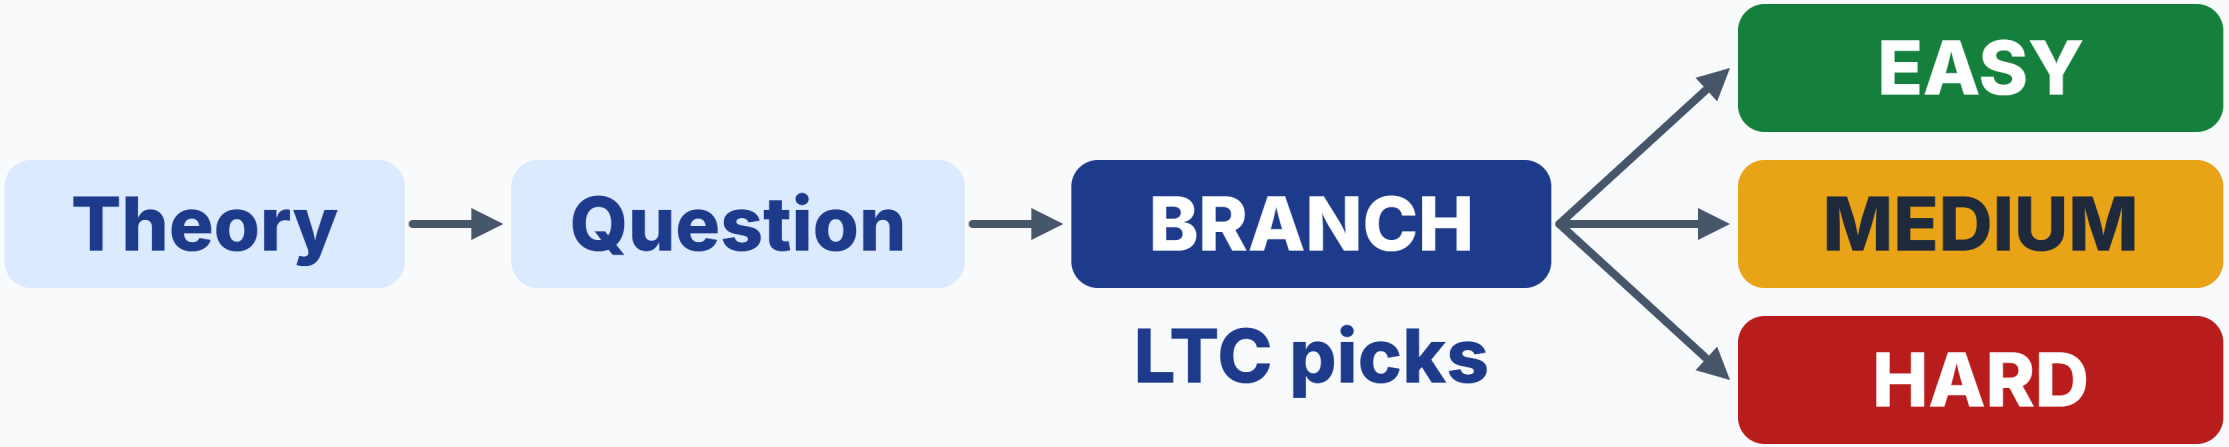
Gap 1

Fixed thresholds ignore psychology

Gap 2

Pacing and feedback are bundled

AdaptEd: the idea



- A lesson is a graph; LTC picks each branch
- Separate feedback: LLM + RAG, teacher-supervised
- Rule fallback · every decision logged

What the network sees: 19 features

9 **Stable profile**
knowledge, mindset, personality, Kolb

5 **Session state**
score, errors in a row, fatigue, time

5 **Item & history**
difficulty, recent accuracy, mastery

The decision is not the score alone

The LTC classifier

$$\frac{dh}{dt} = \frac{-h + \tanh(Wx + Uh + b)}{\tau(h, x)}$$

τ = a “liquid” time constant

3,427
parameters

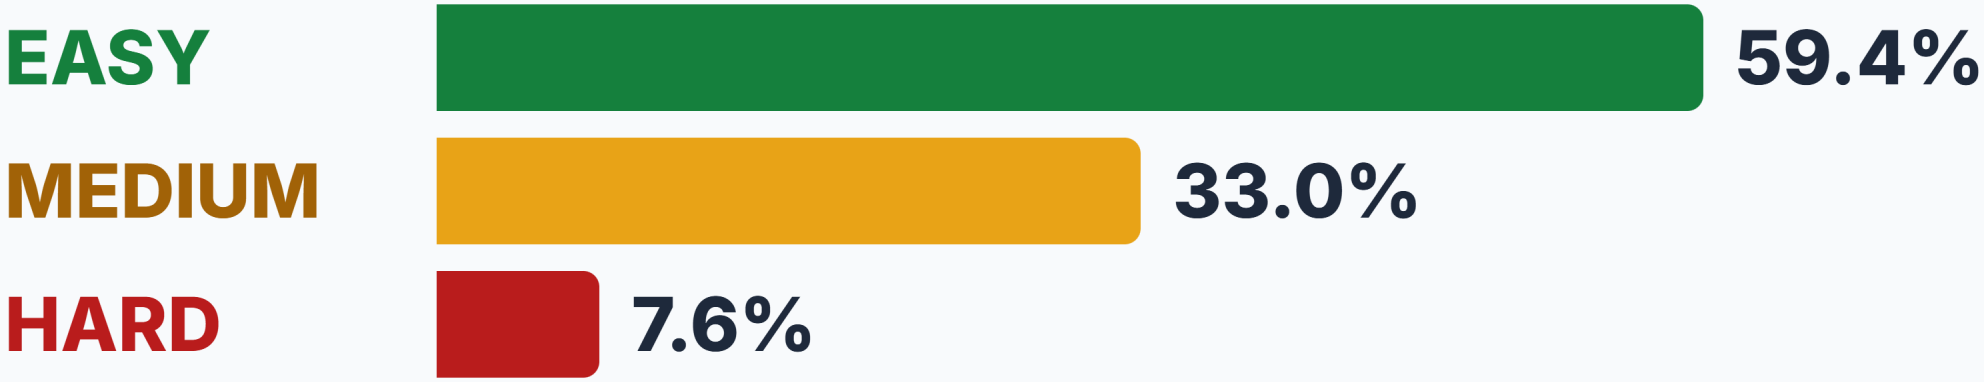
32
neurons

~15 s
training on a CPU

Softmax → **EASY** **MEDIUM** **HARD**

Training data

7,000 **synthetic examples**
5,000 clear + 2,000 boundary cases



Fix: class-balanced loss, model chosen by macro-F1

Results: LTC vs. rules

75.3%
accuracy

0.692
macro-F1

0.636
HARD F1

HARD-class F1



Rules: 85.8% accuracy
expected: labels came from rules

Cost: HARD precision
0.529 vs. 0.828

How it fails, and how sure it is

True ↓ · Predicted →

	EASY	MEDIUM	HARD
EASY	546	78	0
MEDIUM	109	182	56
HARD	0	16	63

All errors: one level apart

8/10

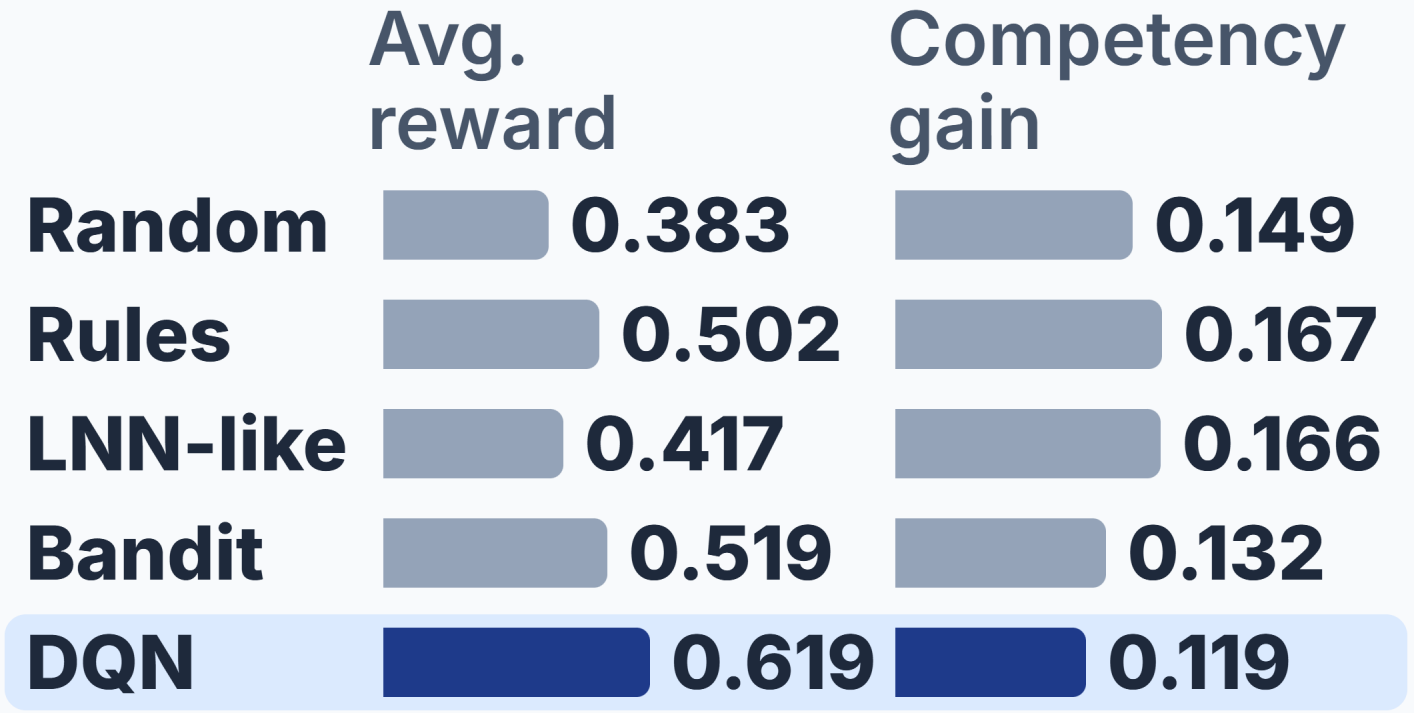
agree with the rules

0.60, 0.61

lowest confidence = the 2 disagreements

Small sample: illustrative only

5 policies, 25,000 episodes each



DQN

highest reward, lowest gain

picks EASY in >92% of steps

= reward hacking

Caveat: “LNN-like” is a heuristic, not our LTC

Takeaways

1. Tiny net (3.4k params) beats rules on HARD
2. Low confidence could flag teacher review
3. Reward design can be the real bottleneck

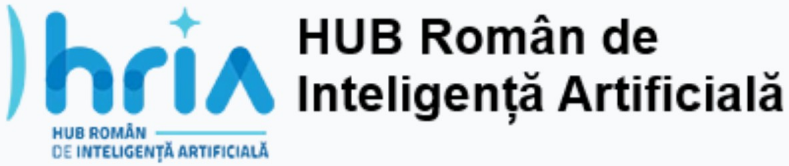
Next

Real students
Trained LTC in simulator
Feedforward baseline

Limitation

All results are on synthetic data

Acknowledgment





Funding

HRIA project
MySMIS 351416

Support

Technical University
of Cluj-Napoca

Thank you! Let's discuss!

Liquid Neural Networks for Adaptive Difficulty Selection in Personalized E-Learning

Sonia Grigor, Constantin Nandra, Dorian Gorgan

Technical University of Cluj-Napoca

sonia.grigor@campus.utcluj.ro