Multi-Objective Test Recommendation for Adaptive Learning

Nassim Bouarour[1], Idir Benouaret [2], Sihem Amer-Yahia[1]
[1]CNRS, Univ. Grenoble Alpes, France
[2]EPITA Research Laboratory (LRE), Le Kremlin-Bicêtre, France

Context

- Upskilling through online platforms (e.g., MOOCs, forums) is rapidly growing but often lacks guidance.
- Learners follow self-directed paths, making it difficult to ensure skill mastery.
- Existing test recommendation systems rarely adapt to individual learning needs or challenges.
- There is a need for personalized, adaptive strategies that drive effective and measurable skill progression.

Motivation

Consider a learner with very basic knowledge in Math who wants to learn mathematical functions.

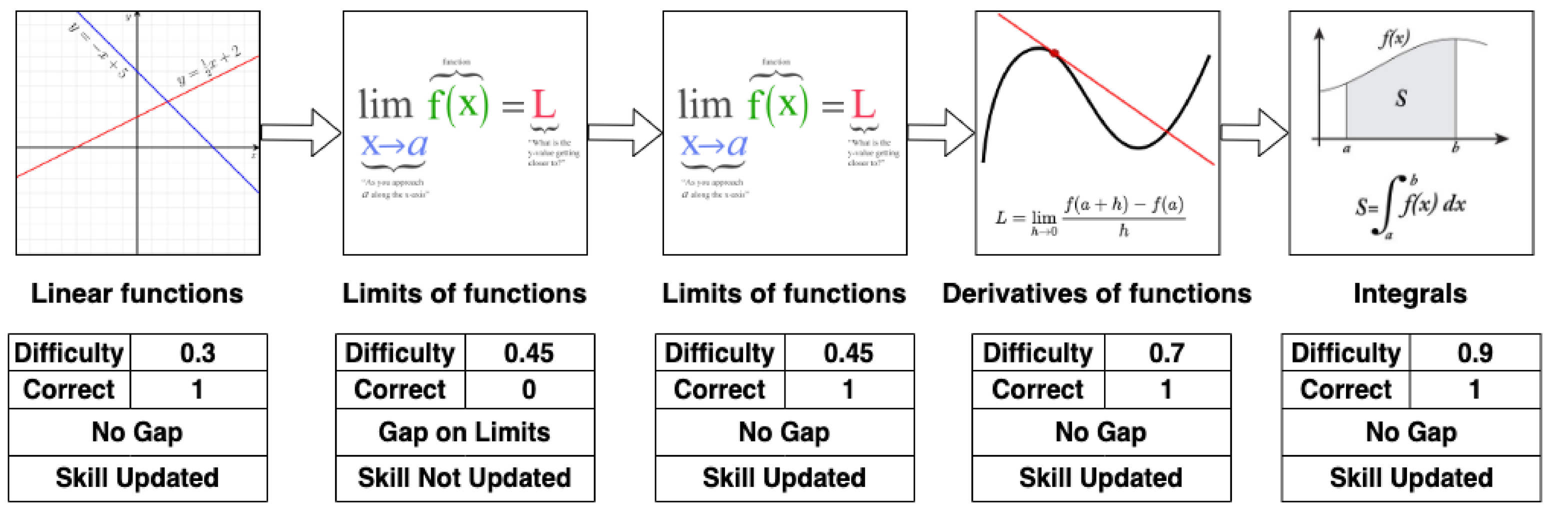

Figure 1: Motivating example

Formalization, Problem statement

We consider a learner $l \in \mathcal{L}$ who follows an iterative learning process for a skill sk. At each step, l completes a set of k tests with different difficulty levels for sk. Each test $t \in \mathcal{T}$ has a fixed difficulty d_t.

Expected performance. It is the expected performance of learner l for a test t. It is based on the similarity of t with successfully completed tests $l.\mathcal{S} \subseteq \mathcal{T}$ by l and is formalized as follows:

$$exPerf(l,t) = sim(t, l.\mathcal{S})$$

Aptitude. It quantifies the difference between a learner's skill value ($l.sk$) and the difficulty level of a test t (d_t). It represents the learner's progression ability for skill sk when assigned tests that are correctly completed. Aptitude is defined as follows:

$$apt(l,t) = d_t - l.sk$$

Gap. It quantifies the distance between the past failed tests of learner l (set $l.\mathcal{F} \subseteq \mathcal{T}$) and the test t and is defined as follows:

$$gap(l,t) = dist(t, l.\mathcal{F})$$

The AdUp problem

To achieve skill mastery, we propose an iterative formulation that solves the following problem:
Given a learner l, with a skill $l.sk$, find a batch $B \subseteq \mathcal{T}$ of k tests to assign to l at iteration i s.t.:

$$\begin{aligned}
&maximize \sum_{t \in B} exPerf(l,t) \\
&maximize \sum_{t \in B} apt(l,t) \\
&minimize \sum_{t \in B} gap(l,t) \\
&\text{subject to } \ |B| = k
\end{aligned}$$

Algorithms

- `MOO`: a Hill Climbing heuristic that finds a subset of the Pareto solutions by optimizing all objectives at once
- We derive multiple variants
- `MAB` each arm corresponds to a subset of dimensions and the reward is the speed of skill progression

Experiments

- **Metrics**: Skill gain, Skill progression, Percentage of learners who attain mastery, Number of iterations to attain mastery

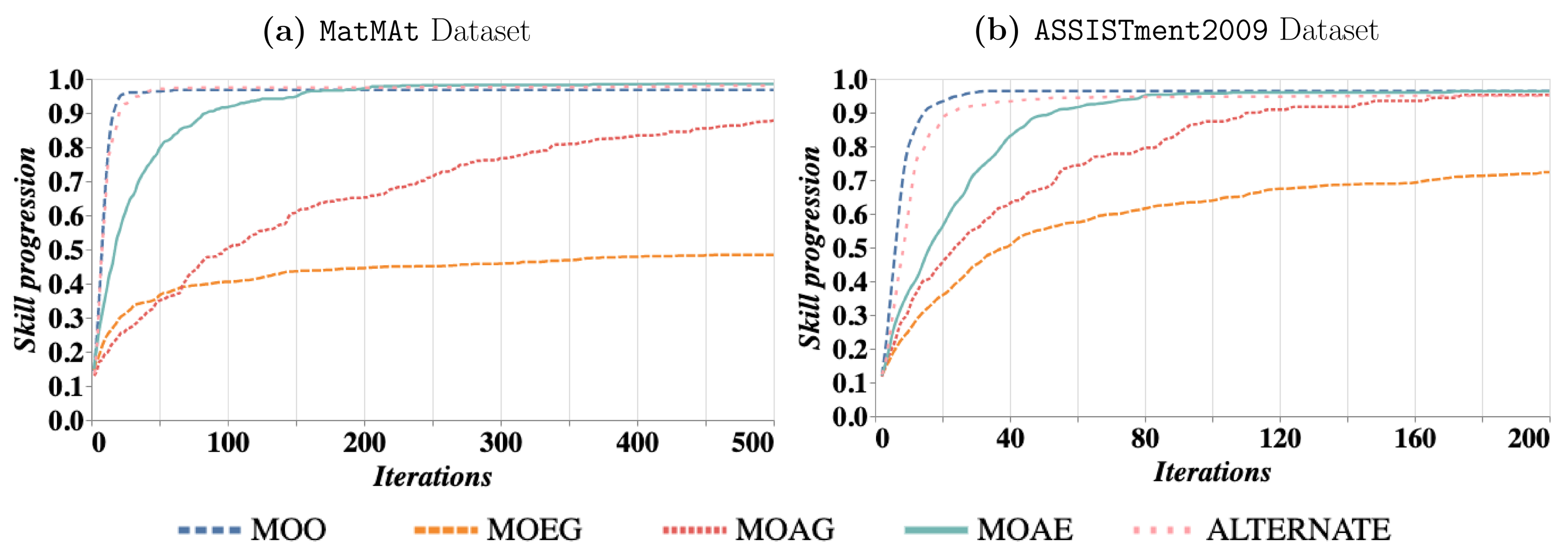

Figure 2: Skill progression as a function of # iterations with a fixed initial skill.

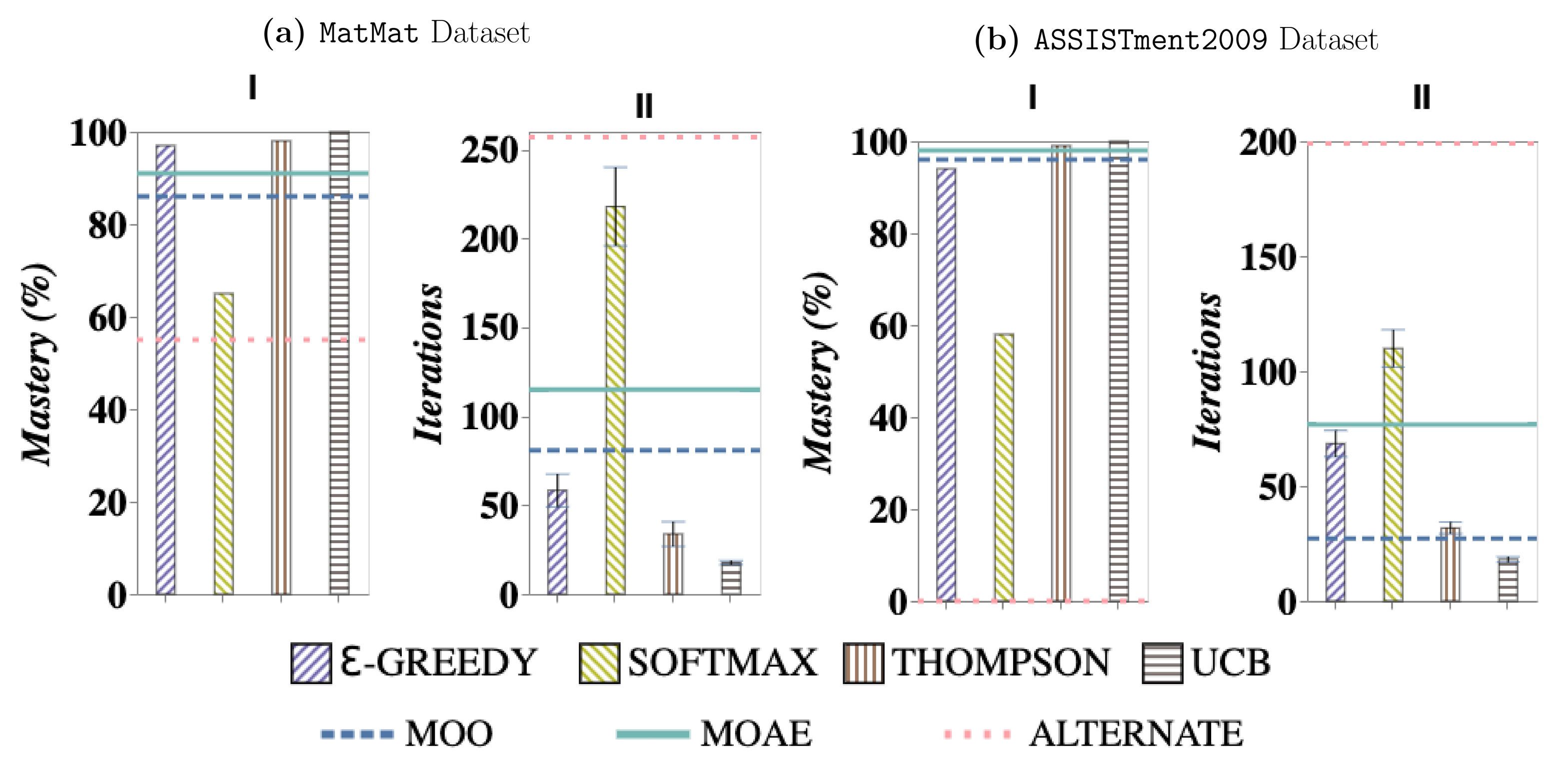

Figure 3: **(I)** Percentage of mastery attained - **(II)** Average number of iterations to attain mastery using `MAB` strategies.

Conclusion

- Novel adaptive upskilling framework, **AdUp**, which recommends personalized test batches to learners by optimizing three key objectives: **expected performance**, **aptitude**, and **skill gap**.
- Two solutions: a **multi-objective optimization (MOO)** method for fixed strategies and a **multi-armed bandit (MAB)** approach that dynamically adapts which objectives to prioritize.
- Through extensive simulations and evaluations on real-world **datasets, our study demonstrates that** combining all three objectives significantly improves **skill gain** and **mastery rate**
- **Adaptive** strategies like MAB outperform static ones by better aligning with **learner progression**. These findings validate the importance of balancing challenge and support in personalized learning systems.

(IA-ÉDU) Intelligence Artificielle générative et ÉDUcation : Enjeux, Défis et Perspectives de Recherche

Code available on GitHub :