Towards AI-Powered Data-Driven Education

Sihem Amer-Yahia
CNRS, Univ. Grenoble Alpes
Grenoble, France
sihem.amer-yahia@univ-grenoble-alpes.fr

ABSTRACT

Educational platforms are increasingly becoming AI-driven. Besides providing a wide range of course filtering options, personalized recommendations of learning material and teachers are driving today's research. While accuracy plays a major role in evaluating those recommendations, many factors must be considered including learner retention, throughput, upskilling ability, equity of learning opportunities, and satisfaction. This creates a tension between learner-centered and platform-centered approaches. I will describe research at the intersection of data-driven recommendations and education theory. This includes multi-objective algorithms that leverage collaboration and affinity in peer learning, studying the impact of learning strategies on platforms and people, and automating the generation of sequences of courses. The paper ends with a discussion of the central role data management systems could play in enabling modern online education.

PVLDB Reference Format:
Sihem Amer-Yahia. Towards AI-Powered Data-Driven Education. PVLDB, 15(12): 3798 - 3806, 2022.
doi:10.14778/3554821.3554900

1 INTRODUCTION

Education is a field of study that deals with the methods and problems of teaching and learning.[1] and constitutes the cornerstone of our societies. The focus in this paper is upskilling, i.e., the process through which one acquires and improves skills. In particular, we are interested in the research questions that arise on online education platforms and how they can benefit from a combination of data-driven approaches, education theory, and artificial intelligence to provide everyone with the ability to advance their knowledge and perfect their skills.

A brief history of education.[28]. Hunter-gatherers had to acquire a vast knowledge of the plants, animals and landscapes on which their survival depended. They had to develop skills in crafting and using tools. They had to be able to take initiative and be creative in finding food and tracking game. 10,000 years later, with the rise of agriculture, and later of industry, people became forced laborers. Land and business owners discovered that they could increase their own wealth by getting other people to work for them. Successful farming and industries required long hours of relatively unskilled, repetitive labor, much of which could be done by children.

Proceedings of the VLDB Endowment, Vol. 15, No. 12 ISSN 2150-8097.
doi:10.14778/3554821.3554900

[1]https://www.merriam-webster.com/dictionary/education

As nations became more centralized, employers saw education as a way to create better workers and leaders saw schooling as means of creating good patriots and soldiers. The most crucial lessons were punctuality, following directions, tolerance for long hours, and a minimal ability to read and write. Education was perceived as community learning. Gurukula [2] is one famous such system and was practised in India in the ancient times. It is a residential schooling system whose origin dates back to around 5000 BC. The students (called shishya) learn from the guru in groups. The essential teachings were in language, science, mathematics through group discussions, and self-learning. Additionally, shishya learn to help the guru in his everyday life, including carrying out mundane daily household chores.

The first degree-granting university is Al-Quarawiyinn in Morrocco. It was founded as a mosque by Fatima al-Fihri in 857–859 and subsequently became one of the leading spiritual and educational centers of the Islamic Golden Age. It was incorporated into Morocco's modern state university system in 1963. The second is Al-Azhar university in Egypt.

The idea and practice of compulsory public education developed gradually in Europe, from the early 16th century to the 19th. By the end of the 17th century, Germany, which was the leader in the development of schooling, had laws in most of its states requiring that children attend school. In China, in 1949, the Communist authorities brought the educational system under national control. To favor China's industrialization, they replaced most studies in humanities and social sciences with engineering. The Chinese Academy of Sciences was set up that same year. Education was reformed and small engineering departments were amalgamated into giant polytechnic institutes such as Tsinghua University. In America, in the mid 17th century, Massachusetts became the first colony to mandate schooling. In Australia, education was made compulsory first in Victoria then Queensland and other regions in the 19th century. In 1872 the Victorian Government passed the Education Act 1872, which set up the colony's public school system that offered free, secular and compulsory education to its children.

Modern education. The same power-assertive methods that had been used to make children work in fields and factories were quite naturally transferred to the classroom. In some schools, children were permitted certain periods of play (recess); but play was not considered to be a vehicle of learning in the classroom. Two major theories are at the basis of today's education systems. *Cognitivism* is a learning theory that emphasizes how information is received, organized, stored and retrieved by the mind. It uses the mind as an information processer, like a computer, and looks beyond observable behavior, viewing learning as internal mental processes. *Constructivism* on the other hand, is a theory in education which

[2]https://en.wikipedia.org/w/index.php?title=Gurukula&redirect=no

posits that learners do not acquire knowledge and understanding by passively perceiving it within a direct process of knowledge transmission, rather they construct new understandings and knowledge through experience and social discourse, integrating new information with what they already know (prior knowledge). These theories lead to emphasizing differences between the need for individual work versus team work, the focus on mentoring approaches, and the confrontation of following a fixed program versus letting students find their way.

Today's work organization necessitates the development of different skills, ones that enable people to be flexible and adaptable in different roles or in different career fields. Specific hard skills and mastery of particular skill sets, with a focus on digital literacy, interaction and collaboration, and managing others, are increasingly in demand. Online education platforms such as Coursera and Udemy, have become a destination for individual upskilling. These platforms play a key role in mediating the success of individuals' education as well as their careers by providing education services. AI has become an enabler of education services (see for instance OntoSIDES, the platform developed in Grenoble to serve medical students and professionals [42]).

2 NEEDS AND CHALLENGES

The need for upskilling in various domains contributed to the rise of online education platforms. Career advancement is considered a right in physical workplaces, but it is still in its infancy online [44]. The next-generation education platforms will need to encompass the ability to learn alone and from others. Several steps must to be taken to enable that: understand how social interactions such as affinity between individuals impact knowledge acquisition in peer learning [22]; understand how teaching strategies that compose learning material of varying difficulty levels, impact skill acquisition and learner's performance and human factors interact with platform goals [37]; provide the ability to set one's goals and automate the composition of teaching material to address those goals [40].

Needs. Learning and upskilling are two human factors that have been extensively studied online [20, 21, 31, 41, 57]. The study of humans factors is a recent trend with various contributions that account for individual factors such as *motivation* [45, 46], *mental stress* [34], and *fatigue and boredom* [12, 30, 50], as well as collaborative factors such as *affinity* and *critical mass* in teams [49]. Human factors can be computed *implicitly* by observing people and their interactions, or *explicitly* by asking them to provide answers via well-established questionnaires.

On peer learning. Online collaboration enables powerful and versatile strategies to improve knowledge of individuals and promote learning. For example, online critiquing communities,[3] social Q&A sites,[4] and crowdsourcing platforms[5] investigate how collaboration can promote knowledge and skill improvement of individuals. It has been shown that the increase in learning one expects from collaboration yields fruitful coordination and higher quality contributions [3, 4]. For instance, in online fan-fiction communities, informal mentoring improves people's writing skills [23]. Coordination between like-minded individuals improves collaboration. This has been verified for crowdsourced text translation [47] where, in addition to skill complementarity, team members' affinity, in terms of age and location, lead to higher quality translation. *In our work on peer learning [22], we propose to explore how affinity between group members improves peer learning, and we address modeling, theoretical, and algorithmic challenges.*

On learning strategies. In physical workplaces, upskilling strategies are regularly implemented and tested [18, 32, 35]. Online, a few studies focused on the role of difficulty and learners' ability to complete tasks in improving their skills [25]. Usually, such approaches require additional human cost to build training material or give feedback to learners. Moreover, there is little understanding of the interplay between achieving high learner performance, a.k.a., quality control and cost reduction, a platform-centric goal, and improving learners' skills and satisfaction, a learner-centric goal. *In our work on studying learning strategies [37], we combine education theory with large-scale user studies to examine the impact of composing teaching material on skill acquisition, quality of contributed content, and learner performance.*

On education pipelines. The current best practices of composing teaching material offer a continuous and consistent process which is mostly done under the guidance of academic advisors. It is needless to say that such a fully manual approach is expensive and inherently not scalable. In contrast, a fully automated approach [5, 14, 15, 24, 43] may require significant historical data to learn personalized models. *Our work on course planning [40] formalizes a sequence generation problem that is sensitive to the ordering and interleaving of items, is personalized and captures progression in task achievement, as well as satisfies a multitude of complex constraints.*

There is a need for new research at the intersection of Data Management, Education Theory, and Machine Learning

Challenges. Research on upskilling and online education strategies requires to launch large-scale user studies to test and validate different solutions. Traditionally, such studies were conducted by gathering subjects in a physical space. Going online calls for a principled approach for sampling subjects, building control and treatment groups, and developing a reproducible experimental protocol. In all our studies, we adopt the same protocol wherein a pre-assessment of subject skills is conducted first by asking subjects to complete a set of tasks and aggregating their skills. That is followed by deploying the actual experiment after which a post-assessment of subjects' skills is done with another set of tasks. The difference between post-assessed and pre-assessed skills constitutes a subject's upskilling.

Gathering data about human subjects is Europe is strictly regulated by the General Data Protection Regulation (GDPR) [6] that requires to inform subjects of the use of their data, where their data is stored and how secure it is, with whom it is shared during the course of the experiment, and for how long it is kept in the system.

[3] https://movielens.org/
[4] http://quora.com/
[5] https://www.figure-eight.com/
[6] https://gdpr-info.eu/

This necessitates to work with data protection lawyers to establish consent forms that are general enough to cover all types of data gathered as well as all the treatments such data will go through. Subjects of our experiments sign such a consent form and have a right to pre-emption, i.e., withdrawal from the experiment, and of asking that their data be permanently deleted.

One of the biggest challenges is to manage human volatility and the evolving nature of humans factors

One major challenge we need to deal with when conducting large scale user studies is to manage human volatility. Since online subjects have a pre-emption right, they may withdraw any time. This necessitates to develop reactive approaches such as incremental optimization to make sure other subjects are not idle for too long.

Another challenge is the many confounding factors due to fluid boundaries between physical and virtual worlds. While user studies are conducted in the virtual world, experiment subjects live in the physical world. Consequently, checking hypotheses via experimental deployment in the virtual world, has to be done carefully.

Finally, the ability to conduct longitudinal observational studies that follow human factors over time is conditioned by the development of back-ends for the management of human factors that will enable the creation of experimental narratives by making the results of user studies persistent and comparable.

3 REVIEW OF SOME WORK

3.1 Peer learning

Peer learning is a form of cooperative learning that relies on explicit/implicit social interactions where one or more individuals act as peer teacher(s), and others experience skill improvement, a.k.a. upskilling. In our work [22], we explored how affinity impacts learning potential in teams. We formalized learning potential and affinity structures and developed novel team formation algorithms with provable theoretical guarantees.

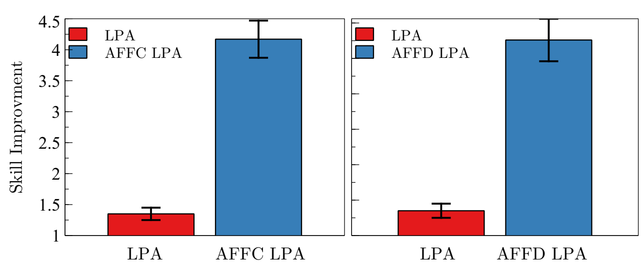

Figure 1: Skill improvement with and w/o affinity in LpA

Team formation in online communities [7, 8, 36, 47, 49] is often stated as: given a set of individuals and tasks, form a set of groups for the tasks that optimize some aggregated utility subject to constraints such as team size, maximum workload etc. Utility can be aggregated in different ways: the sum of individual skills, their product, etc [8]. Team formation is combinatorial in nature and proposed algorithms solve the problem under different constraints and utility definitions (e.g., [36]). Unlike these problems, we study how to form teams with the goal of maximizing peer learning under different affinities.

Our first contribution is to present principled models to formalize peer learning and affinity structures. The learning potential of a peer from a more skilled peer is naturally defined as the skill difference between the latter and the former [3, 4]. We use that to formulate two common learning models (see Figure 2): LpA where each member learns from all higher skilled ones, and LpD where the least skilled member (resp., the most skilled) learns from (resp. teaches to) all others.

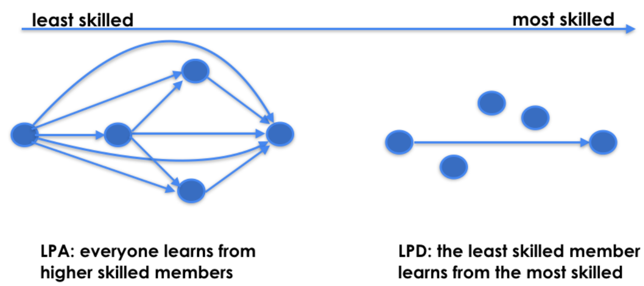

Figure 2: Two formulations of learning potential

Affinity, on the other hand, depends on the application and can be expressed using common socio-demographic attributes or more generally, using models that capture psychological traits such as the Myers-Briggs test [11]. We study our two learning models in conjunction with two common affinity structures (see Figure 3): AffD where group affinity is a function of all pairwise affinities between its members, and AffC where it depends on affinities with a moderator.

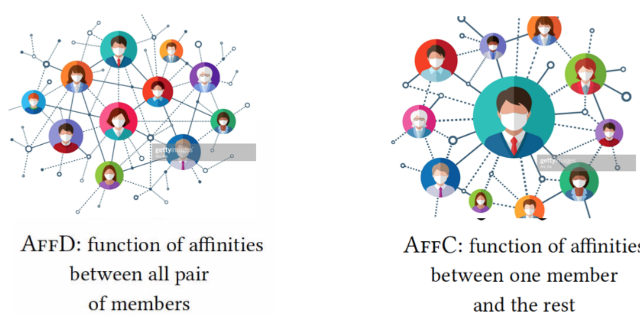

Figure 3: Two formulations of affinity

Different combinations of learning models and affinity structures capture different real-world cases. For instance, for fact-checking tasks, each member learns from those who better know the facts, and the two least collaborative individuals must get along. For fact-learning tasks, each member learns from those who better know the facts, and all members must get along with the moderator. We investigate these two kinds of tasks in our experiments and observe that AffD LpA best represents groups whose goal is to check facts, while AffC LpA is more appropriate for fact-learning.

Our second contribution is to study the formalized models systematically and present our theoretical findings. Our original problem formulations are bi-objective optimizations, with the goals to build k equi-sized groups over a set of n members that maximize both learning potential and affinity. Interestingly, we prove that both learning potential variants, LpD and LpA, can be solved in polynomial time by sorting individuals on their skills, however, the problems become NP-hard when affinity and group size constraints are considered. Therefore, our solution first finds k groups that yield the highest possible learning potential value and then transforms our two-objective problem into *a constrained optimization* that looks for k groups that optimize affinity, with that learning potential value as a constraint.

Problem	Algo.	Approx.	Time
(AffC LpD)	GrAffC-LpD	exact LpD, 3 AffC	$O(klogn + nlogk)$
(AffC LpA)	GrAffC-LpA	exact LpD, 3 AffD	$O(nlogn)$
(AffD LpD)	GrAffD-LpD	exact LpA, 6 AffC	$O(klogn + nlogk)$
(AffD LpA)	GrAffD-LpA	exact LpA, 6 AffD	$O(nlogn)$

Figure 4: Our technical contributions on peer learning with affinities.

Our third contribution is algorithmic. We present a suite of scalable algorithms that form groups to maximize learning potential and optimize affinity within *constant approximation factors*. To attain their approximation guarantees, these algorithms assume that affinity satisfies triangle inequality [36]. Many similarity/distance measures such as Jaccard distance and edit distance are known to satisfy metric properties and these properties are usually assumed to design algorithms with guarantees [36]. Our technical contributions are summarized in Figure 4.

3.2 Learning strategies

In our work on learning strategies [37], we study how composing learning material of various difficulty levels impacts the inherent skill improvement of humans and their overall performance. We focus on a common class of tasks referred to as “Knowledge and Comprehension tasks” in Bloom’s taxonomy of educational objectives [10, 33] such as image classification, labeling, editing grammar&spelling mistakes, and speech transcription.

There are two education theories underlying our framework. First, Zone of Proximal Development (ZPD) [56] is a well-known theory that defines three zones of tasks with different skill improvements; (1) A learnable zone that contains tasks a person can learn how to complete when assisted by a teacher or peer with a higher skill set, (2) a flow/comfort zone of tasks that are easy and can be completed with no help, and (3) a frustration zone of tasks that a learner cannot complete even with help. Second, the Flow theory [16] states that people are able to immerse themselves in doing things whose challenge matches their skills. Figure 5 integrates the two theories and illustrates their relationship with respect to the task challenge, the subject skill, and the affect state [9, 53]. In [9], the authors claim that to improve skills, the tasks should be either in the flow/comfort zone, or in the learnable zone on the condition that there is some “scaffolding” to help subjects complete tasks that are a bit more challenging for them. This results in skill

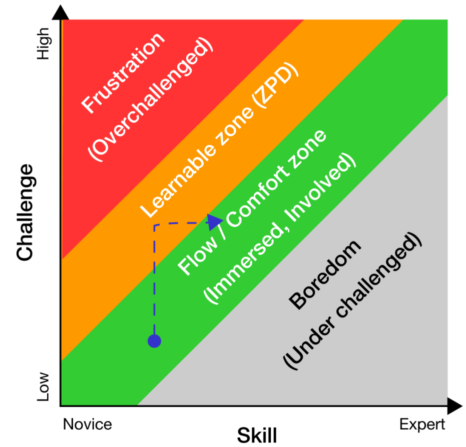

Figure 5: Zone of Proximal Flow [9], which combines the results of Zone of Proximal Development and Flow Theory. Scaffolding tasks helps learners improve their skills by completing more challenging tasks (the dotted line).

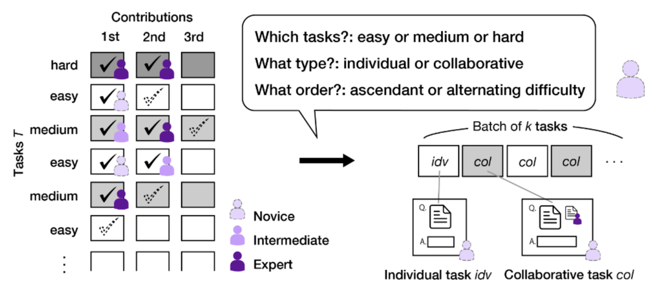

Figure 6: Learning strategies: Given a learner r and a set of tasks T, generate a sequence of k tasks to maximize learning.

improvement (the dotted line). Our formalization builds on that and defines the learning potential for both individual tasks (mainly in the flow/comfort zone) and collaborative tasks (mainly in the learnable zone).

Our question is illustrated in Figure 6. We have a set T of tasks of varying difficulty levels, each task receives N (=3 in the figure) contributions. At each iteration, some tasks have already been completed by some subjects.

Given a learner r and the batches of tasks completed by r up to iteration i: $B_1 \ldots B_i$, find a batch B of at most k tasks to assign to r at iteration $i+1$ such that it maximizes the learner’s upskilling. Here, learning potential is the maximum possible improvement in w’s skill. We assume that subject skill and task difficulty are uni-dimensional and that the skill of a subject either remains the same or increases as time passes [38, 55].

Our challenges are: How to choose an appropriate batch of k tasks where a subject can see previous higher-skilled subjects’ contributions? How to order the k tasks so that the subject’s skill improvement is maximized? How to reconcile subject-centric and platform-centric goals?

Our problem can be solved as a knapsack problem and top-k search can be computed linearly in the number of tasks.

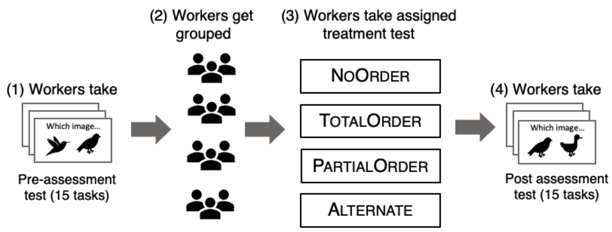

Figure 7: Experimental protocol for deploying our learning strategies for individual tasks

We devise learning strategies which build on two ideas: (1) task ordering, and (2) interleaving individual and collaborative tasks. We study their impact on subjects' performance and skills. Previous work found that both task ordering and task types impact contribution quality and completion time [13, 17, 19]. That is the basis for designing our four task orderings: NoOrder, a baseline where tasks are in no particular order; TotalOrder, where tasks are presented in increasing difficulty level, PartialOrder, a variation of TotalOrder, where tasks are grouped according to their difficulty and groups presented in increasing difficulty; and Alternate, that groups tasks and presents them in alternating difficulty levels.

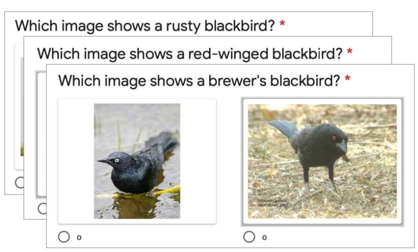

Figure 8: Individual tasks

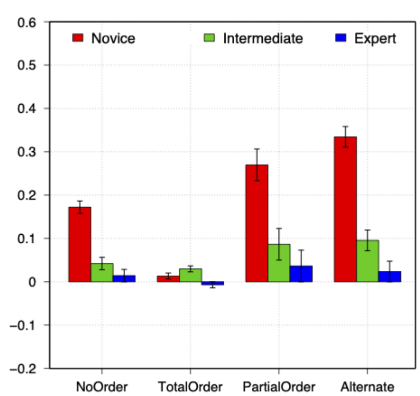

Figure 9: Results of learning strategies for individual tasks

Our subjects go through the following protocol (see Figure 7). A subject takes a pre-assessment test following which the subject is assigned to a treatment group and takes the corresponding treatment, i.e., task ordering. After completing the treatment test, the subject takes the post-assessment test, which is similar to the pre-test but with different questions. Upskilling is computed as a difference between post-assessed and pre-assessed skills.

We empirically studied the impact of learning strategies applied to individual tasks. The task is simply to identify the specified bird type given a pair of bird images (see Figure 8). We conducted two variants of this experiment: one with 12 tasks and another with 120 tasks. Results are reported in Figure 9. We observe that all treatments are effective in helping learners improve their skill. In particular, Alternate yielded the highest average skill improvement in both 12-task and 120-task experiments. Further, it is the best strategy for both novice and intermediate learners while there is a ceiling effect for expert learners.

We further examined the learners' upskilling and performance in a series of interleaved collaborative tasks (CTs) and individual tasks of text editing (See Figure 10). We asked subjects to correct English spelling and grammar errors. In the collaborative version, they can see answers of higher-skilled subjects and have the option to edit the current answer. We asked subjects to take a pre-assessment test to measure their English skills and based on the test result, they were classified as novice-intermediate or expert. Novice-intermediate subjects then work using either the CTs only strategy or the interleaved strategy. Lastly, novice-intermediate subjects take a post-test. Our experiments confirmed the observations we had in the individual task experiment. Skill improvement is significantly higher in the interleaved case compared to the case where subjects had CTs only. The higher skill improvement in the interleaved case may be attributed to the fact that individual tasks are similar to CTs, which may have contributed to the learning of the subjects. Moreover, in the case of CTs only, since there are already answers from expert subjects, the novice-intermediate subjects may have become under-challenged, resulting in a lower skill improvement. Throughput is also observed to be higher in the interleaved case. We can conjecture that as skills improve, subjects become more proficient and faster.

Our experiments show, with statistical significance, that the learning strategies are effective in helping subjects improve their skills. More specifically, Alternate yields the highest average skill improvement for individual tasks, and subjects produce the highest quality contributions, best task throughput, and highest skill improvement, when collaborative and individual tasks are interleaved. We can hence conclude by saying that hypotheses verified in physical workplaces also apply in virtual marketplaces: the alternation of task difficulty yields the highest upskilling and throughput.

3.3 Education pipelines

In our work on automating education pipelines [40], we propose a computational framework to automate the generation of course plans that is applicable to a variety of domains. Scenarios where in-person education is rare and costly and platforms that need to scale up the process to thousands of items such as MOOCs [51] are ideal for our problem.

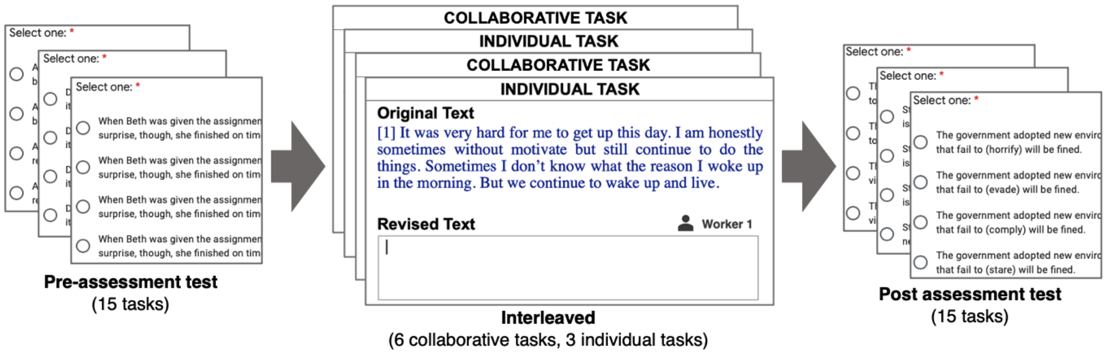

Figure 10: Deployment of interleaved individual and collaborative tasks

Consider an aspiring youngster wanting to jump-start her career as a data scientist right after her B.S. in Computer Science, or a seasoned IT analyst with years of experience in industry wanting to join the bandwagon of data science to change her career focus. For both individuals, designing a course plan is a complex and intellectually demanding task with the goal of managing their upskilling goal, and satisfying requirements that are compatible with their experience and background.

Our example calls out the following requirements in course planning - (1) Satisfying Hard Constraints: Plans must match these constraints (e.g., # core vs. # elective courses, as well as prerequisite requirements). (2) Maximizing Soft Constraints: These are of two kinds: (a) Designed plans must maximize the coverage of the topics/themes a learner wishes to acquire (e.g., recommend courses on clustering and neural networks); (b) Recommended sequences must adhere as much as possible to a "template" provided by a domain expert to reflect ideal permutations of core and elective courses.

Our first contribution is to formalize a constrained sequence generation problem. We model a Constrained Markov Decision Process [6] where a state is a course, an action generates a transition that adds one or more course, and a "reward" is associated with every transition to quantify how well the action satisfies the hard constraints, and maximizes the soft constraints. Designing a reward function that captures all these nuances is a complex and intellectually demanding data science task.

Our second contribution is to present a computational framework, that is inspired by Constrained Reinforcement Learning (C-RL) [2, 27], specifically Weighted RL [26], but non-trivially adapts it to handle multiple hard and soft constraints. Essentially, we propose a weighted reward function to transform the Markov process to an unconstrained process that captures multiple hard constraints as well as maximizes the actual value by maximizing the soft constraints. We prove that our designed reward function satisfies all hard constraints. We adapt the popular model-free on-policy algorithm SARSA [52] for updating the Q values of the states, that is known to converge faster and with fewer errors [48].

Our third contribution is an extensive evaluation using two real datasets to plan courses for 4 different sought after degree programs. Our results convincingly demonstrate that: (a) Our algorithm generates course plans that are comparable in quality to handcrafted ones, and are superior to fully automated sequence-aware recommendations (e.g., *OMEGA* [54]) and to next-step recommendation in *EDA* [39]; (b) based on user studies involving 25 data science computational track (DS-CT) major students, our course plans achieve highly comparable satisfaction scores w.r.t. handcrafted gold standards designed by domain experts; (c) the policy learned by our solution, `RL-Planner`, for the M.S. DS-CT is transferable to a different degree program in M.S. Computer Science inside the same university and vice versa; (d) our algorithm is robust to the different parameters, takes reasonable time for learning the policy, and can therefore make interactive recommendations. More specifically, our experiments address the following questions:

Q1. How well `RL-Planner` performs in comparison to baselines?
Q2. How do end users (students or travelers) compare recommendations by `RL-Planner` to gold standards?
Q3. How effective is `RL-Planner` for transfer learning?
Q4. How robust is `RL-Planner` w.r.t. different parameters?
Q5. How scalable is `RL-Planner`?

To answer Q1 and Q4, we present average scores over 10 runs. For Q2, we run a user study and measure user satisfaction in a $1-5$ scale. We describe two case studies to answer Q4. The score of each recommendation is computed as the highest distance with each ideal composition $I \in IT$ to capture satisfaction of soft constraints. Finally, we study running time to answer Q5.

We consider datasets extracted from the NJIT (Univ-1) and Stanford (Univ-2) websites. Univ-1 (NJIT) contains 1216 courses comprising 126 degree programs through 6 professional schools and colleges. We focus on 3 M.S. programs: Data Science-Computational Track (DS-CT), Cybersecurity, and Computer Science (CS). The hard constraints consider the number of core and elective courses while satisfying the gap between a course and its prerequisites. Univ-2 (Stanford) contains 3742 courses for 4 different departments related to data science. Each course has a title, department number, department code, course description, prerequisites, minimum and maximum number of required units. We focus on the M.S. Data Science (DS) program. The hard constraints are designed considering the number of units constraints in the following 6 sub-disciplines while satisfying prerequisites gaps: a. Mathematical and Statistical Foundations; b. Experimentation; c. Scientific Computing (includes software development and large-scale computing); d. Applied Machine Learning and Data Science; e. Practical Component; f. Elective course in the data science. To form topic vectors, we extract nouns from course names and remove stopwords. In Univ-1, we get 60, 61, and 100 distinct topics for DS-CT, Cybersecurity, and CS. We obtain 73 topics from Univ-2.

We have two baselines. Manual Gold Standard, a handcrafted sequence of courses designed by academic advisors for the relevant degree programs at Univ-1. For Univ-2, we obtain the gold standard from the website of the degree program. The gold standard scores are 10 for Univ-1 and 15 for Univ-2, since the ideal course plans consist of 10 and 15 courses, respectively. Automated Solutions, one that performs sequence mining, and the other that adapts exploratory data analysis (*EDA*), cannot be adapted to transfer learning.

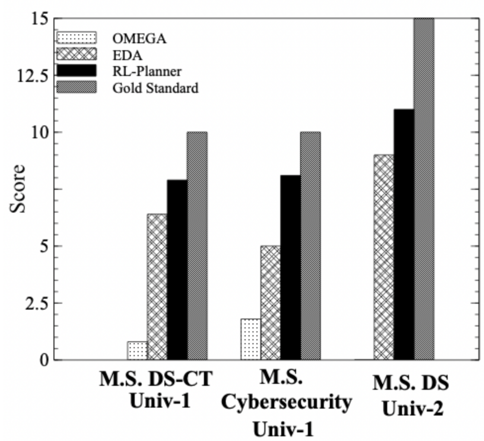

Figure 11: `RL-Planner` **vs *OMEGA, EDA,* and Gold Standard**

Our results demonstrate that: (a) Existing fully automated approaches are not capable to adapt to sequence recommendations with complex constraints. Both *OMEGA* [54] and *EDA* are unable to generate course plans that satisfy hard constraints most of the time, `RL-Planner` generates high quality course plans that are comparable to handcrafted gold standards; (b) Based on user studies involving 25 data science computational track (DS-CT) major students, `RL-Planner` is highly comparable w.r.t. handcrafted gold standards. `RL-Planner` gets 3.39 user satisfaction score on average out of 5 compared to 3.74 for gold standards. (c) `RL-Planner` is effective in transferring policies between different Master's programs, whereas, the fully automated baselines cannot; (d) `RL-Planner` takes reasonable time for learning the policy, and is capable to make interactive recommendations in real time.

We compare the plans generated by `RL-Planner` to baselines *OMEGA*, *EDA* and to the fully manual gold standard. Figure 11 presents the average scores where a score is computed as a distance from the interleaving template to reflect how well the soft constraint is satisfied. We observe that `RL-Planner` generates plans that are higher in score than the fully automated baselines for all cases while being very close to the gold standard. This confirms the suitability of reinforcement learning for training education pipelines under constraints.

4 OPEN RESEARCH QUESTIONS

We are now ready to define AI-powered data-driven education:

AI-powered data-driven education is the use of data and AI algorithms to devise learning strategies and study their impact on humans, and on platforms

We conclude with a discussion of three open questions.

OQ1: Peer recommendation. This question relates to finding peers who serve as teachers and help one upskill. This is different from people recommendation [1] as in dating systems where the recommendation is reciprocal. Peer recommendation is asymmetric and concerns a learner and one or multiple teachers. The exact identity of teachers need not be fixed. Rather, it is the profile of recommended teachers that needs to be determined and how they complement a learner's profile. This could be powered by a profile management database that contains profile descriptions rather than specific individuals. Once a profile is recommended to a learner, it could be used to find actual subjects on different platforms that best fit that profile. This will help handle human volatility, i.e., the fact that some subjects may or not actually commit to the experiment. By managing profiles, the system will be able to quickly identify a replacement, especially if the task at hand is collaborative and involves other learners and teachers, as in team formation in Section 3.1, who cannot be left idle for too long.

OQ2: Compose humans and learning material. This question relates to the ability to combine teaching material, in the form of individual or collaborative tasks, with humans. This composition must incorporate learning strategies as described in Section 3.2. The challenge here is to express optimization problems whose solution is an education pipeline that implements a strategy that best combines tasks of various difficulty levels. Reinforcement learning appears to be a relevant approach here where a state is a task/learning material, an action is the addition of that material to the pipeline, and a reward captures a learning strategy. The deployment of these strategies in practice needs to account for other human factors, in particular, fatigue and boredom [29]. Capturing those factors while people are consuming learning material, and feeding them into the composition process is an open research question.

OQ3: Learner feedback. The last question pertains to providing learners with the ability to express their goals in a declarative manner as shown in Section 3.3. This must be complemented with the design and deployment of longitudinal observational studies of human factors over time to observe individuals as they consume learning material, quantify their performance and growth, and adapt optimization goals and solutions accordingly. Leveraging ML approaches will enable learning profiles dynamically. The question of revisiting optimization goals would require to rethink the way algorithms adapt to this dynamicity by learning the relationship between profiles and needs over time. This has the potential to impact new domains that go beyond education thereby widening the scope of "AI-powered and data-driven" research.

ACKNOWLEDGMENTS

This work has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 863410.

Special thanks go to Marie-Christine Rousset, Divesh Srivastava, Shady Elbassuoni, for their feedback on earlier versions of my talk.

REFERENCES

[1] Himan Abdollahpouri, Gediminas Adomavicius, Robin Burke, Ido Guy, Dietmar Jannach, Toshihiro Kamishima, Jan Krasnodebski, and Luiz Augusto Pizzato. 2020. Multistakeholder recommendation: Survey and research directions. *User Model. User Adapt. Interact.* 30, 1 (2020), 127–158.

[2] Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. 2017. Constrained policy optimization. In *ICML*.

[3] Rakesh Agrawal, Behzad Golshan, and Evimaria Terzi. 2014. Grouping students in educational settings. In *SIGKDD*.

[4] Rakesh Agrawal, Sharad Nandanwar, and Narasimha Murty Musti. 2017. Grouping Students for Maximizing Learning from Peers. In *EDM*.

[5] Sunita B Aher and LMRJ Lobo. 2013. Combination of machine learning algorithms for recommendation of courses in E-Learning System based on historical data. *Knowledge-Based Systems* 51 (2013), 1–14.

[6] Eitan Altman. 1999. *Constrained Markov decision processes*. Vol. 7. CRC Press.

[7] Aris Anagnostopoulos et al. 2010. Power in unity: forming teams in large-scale community systems. In *CIKM*.

[8] Aris Anagnostopoulos et al. 2012. Online team formation in social networks. In *WWW*.

[9] Ashok R Basawapatna, Alexander Repenning, Kyu Han Koh, and Hilarie Nickerson. 2013. The zones of proximal flow: guiding students through a space of computational thinking skills and challenges. In *Proceedings of the ninth annual international ACM conference on International computing education research*. 67–74.

[10] Benjamin S Bloom. 1956. Taxonomy of educational objectives. Vol. 1: Cognitive domain. *New York: McKay* (1956), 20–24.

[11] Katharine C Briggs. 1976. *Myers-Briggs type indicator*. Consulting Psychologists Press Palo Alto, CA.

[12] Alec Burmania, Srinivas Parthasarathy, and Carlos Busso. 2015. Increasing the reliability of crowdsourcing evaluations using online quality assessment. *IEEE Transactions on Affective Computing* 7, 4 (2015), 374–388.

[13] Carrie J Cai, Shamsi T Iqbal, and Jaime Teevan. 2016. Chain reactions: The impact of order on microtask chains. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 3143–3154.

[14] Munmun De Choudhury et al. 2010. Automatic construction of travel itineraries using social breadcrumbs. In *Hypertext*. 35–44.

[15] Ko-Kang Chu et al. 2003. Designing a course recommendation system on web based on the students' course selection records. In *EdMedia+ Innovate Learning*. AACE, 14–21.

[16] Mihaly Csikszentmihalyi. 1975. *Beyond boredom and anxiety: The experience of play in work and games*. Jossey-Bass.

[17] Peng Dai, Jeffrey M Rzeszotarski, Praveen Paritosh, and Ed H Chi. 2015. And now for something completely different: Improving crowdsourcing workflows with micro-diversions. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*. ACM, 628–638.

[18] Leo J De Vin, Lasse Jacobsson, JanErik Odhe, and Anders Wickberg. 2017. Lean Production Training for the Manufacturing Industry: Experiences from Karlstad Lean Factory. *Procedia Manufacturing* 11 (2017), 1019–1026.

[19] Djellel Eddine Difallah, Gianluca Demartini, and Philippe Cudré-Mauroux. 2016. Scheduling human intelligence tasks in multi-tenant crowd-powered systems. In *Proceedings of the 25th international conference on World Wide Web*. 855–865.

[20] Mira Dontcheva, Robert R Morris, Joel R Brandt, and Elizabeth M Gerber. 2014. Combining crowdsourcing and learning to improve engagement and performance. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 3379–3388.

[21] Shayan Doroudi, Ece Kamar, and Emma Brunskill. 2019. Not Everyone Writes Good Examples but Good Examples Can Come from Anywhere. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 7. 12–21.

[22] Mohammadreza Esfandiari, Dong Wei, Sihem Amer-Yahia, and Senjuti Basu Roy. 2019. Optimizing Peer Learning in Online Groups with Affinities. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1216–1226.

[23] Sarah Evans et al. 2017. More Than Peer Production: Fanfiction Communities as Sites of Distributed Mentoring. In *CSCW*.

[24] Zachary Friggstad et al. 2018. Orienteering Algorithms for Generating Travel Itineraries. In *WSDM*. 180–188.

[25] Ujwal Gadiraju and Stefan Dietze. 2017. Improving learning through achievement priming in crowdsourced information finding microtasks. In *Proceedings of the Seventh International Learning Analytics & Knowledge Conference*. ACM, 105–114.

[26] Peter Geibel. 2006. Reinforcement learning for MDPs with constraints. In *European Conference on Machine Learning*. Springer, 646–653.

[27] Peter Geibel and Fritz Wysotzki. 2005. Risk-sensitive reinforcement learning applied to control under constraints. *JAIR* 24 (2005), 81–108.

[28] P. Gray. 2013. *Free to Learn: Why Unleashing the Instinct to Play Will Make Our Children Happier, More Self-Reliant, and Better Students for Life*. Basic Books. https://books.google.fr/books?id=2IgkAIMG3qgC

[29] Kenji Hata, Ranjay Krishna, Li Fei-Fei, and Michael S Bernstein. 2017. A glimpse far into the future: Understanding long-term crowd worker quality. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. ACM, 889–901.

[30] Kazushi Ikeda and Keiichiro Hoashi. 2017. Crowdsourcing Go: Effect of worker situation on mobile crowdsourcing performance. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 1142–1153.

[31] Juho Kim. 2015. *Learnersourcing: improving learning with collective learner activity*. Ph.D. Dissertation. Massachusetts Institute of Technology.

[32] Martijn Koops and Martijn Hoevenaar. 2013. Conceptual change during a serious game: Using a Lemniscate model to compare strategies in a physics game. *Simulation & Gaming* 44, 4 (2013), 544–561.

[33] David R Krathwohl and Lorin W Anderson. 2009. *A taxonomy for learning, teaching, and assessing: A revision of Bloom's taxonomy of educational objectives*. Longman.

[34] Katsumi Kumai, Masaki Matsubara, Yuhki Shiraishi, Daisuke Wakatsuki, Jianwei Zhang, Takeaki Shionome, Hiroyuki Kitagawa, and Atsuyuki Morishima. 2018. Skill-and-stress-aware assignment of crowd-worker groups to task streams. In *Sixth AAAI Conference on Human Computation and Crowdsourcing*.

[35] Susanne P Lajoie and Alan Lesgold. 1989. Apprenticeship training in the workplace: Computer-coached practice environment as a new form of apprenticeship. *Machine-mediated learning* 3, 1 (1989), 7–28.

[36] Theodoros Lappas, Kun Liu, and Evimaria Terzi. 2009. Finding a team of experts in social networks. In *SIGKDD*.

[37] Masaki Matsubara, Ria Mae Borromeo, Sihem Amer-Yahia, and Atsuyuki Morishima. 2021. Task Assignment Strategies for Crowd Worker Ability Improvement. *Proc. ACM Hum. Comput. Interact.* 5, CSCW2 (2021), 1–20.

[38] Julian John McAuley and Jure Leskovec. 2013. From amateurs to connoisseurs: modeling the evolution of user expertise through online reviews. In *Proceedings of the 22nd international conference on World Wide Web*. 897–908.

[39] Tova Milo and Amit Somech. 2018. Next-step suggestions for modern interactive data analysis platforms. In *SIGKDD*. 576–585.

[40] Sepideh Nikookar, Pars Sakharkar, Baldjinder Smagh, Sihem Amer-Yahia, and Senjuti Basu Roy. 2022. Guided Task Planning Under Complex Constraints. (2022).

[41] Nilesh Padhariya and Kshama Raichura. 2014. Crowdlearning: An incentive-based learning platform for crowd. In *2014 Seventh International Conference on Contemporary Computing (IC3)*. IEEE, 44–49.

[42] Olivier Palombi, Fabrice Jouanot, Nafissetou Nziengam, Behrooz Omidvar-Tehrani, Marie-Christine Rousset, and Adam Sanchez. 2019. OntoSIDES: Ontology-based student progress monitoring on the national evaluation system of French Medical Schools. *Artif. Intell. Medicine* 96 (2019), 59–67.

[43] Aditya Parameswaran et al. 2011. Recommendation systems with complex constraints: A course recommendation perspective. *TOIS* (2011).

[44] Participants of Shonan Future-of-Work Workshop. 2020. Data Management Research and More: Making AI Machines Work for Humans in FoW. In *ACM SIGMOD Record (to appear)*.

[45] Julien Pilourdault, Sihem Amer-Yahia, Dongwon Lee, and Senjuti Roy. 2017. Motivation-aware task assignment in crowdsourcing. In *EDBT*.

[46] Julien Pilourdault, Sihem Amer-Yahia, Senjuti Basu Roy, and Dongwon Lee. 2018. Task relevance and diversity as worker motivation in crowdsourcing. In *2018 IEEE 34th International Conference on Data Engineering (ICDE)*. IEEE, 365–376.

[47] Habibur Rahman et al. 2015. Task assignment optimization in collaborative crowdsourcing. In *ICDM*.

[48] Sudharsan Ravichandiran. 2018. *Hands-on Reinforcement Learning with Python: Master Reinforcement and Deep Reinforcement Learning Using OpenAI Gym and TensorFlow*. Packt Publishing Ltd.

[49] Senjuti Basu Roy, Ioanna Lykourentzou, Saravanan Thirumuruganathan, Sihem Amer-Yahia, and Gautam Das. 2015. Task assignment optimization in knowledge-intensive crowdsourcing. *VLDB J.* 24, 4 (2015), 467–491.

[50] Jeffrey M Rzeszotarski, Ed Chi, Praveen Paritosh, and Peng Dai. 2013. Inserting micro-breaks into crowdsourcing workflows. In *First AAAI Conference on Human Computation and Crowdsourcing*.

[51] Charles Severance. 2012. Teaching the world: Daphne koller and coursera. *Computer* 45, 8 (2012), 8–9.

[52] Richard S Sutton and Andrew G Barto. 2018. *Reinforcement learning: An introduction*. MIT press.

[53] Mohammad Taheri, Nasser Sherkat, Nick Shopland, Dorothea Tsatsou, Enrique Hortal Nicholas Vretos, Christos Athanasiadis, and Penny Standen. 2017. Adaptation and Personalization principles based on MaThiSiS findings. In *Public report on Managing Affective-learning THrough Intelligent atoms and Smart InteractionS project*.
[54] Sebastian Tschiatschek, Adish Singla, and Andreas Krause. 2017. Selecting Sequences of Items via Submodular Maximization.. In *AAAI*. 2667–2673.
[55] Kazutoshi Umemoto, Tova Milo, and Masaru Kitsuregawa. 2020. Toward Recommendation for Upskilling: Modeling Skill Improvement and Item Difficulty in Action Sequences. In *2020 IEEE 36th International Conference on Data Engineering (ICDE)*. IEEE, 169–180.
[56] Lev Vygotsky. 1987. Zone of proximal development. *Mind in society: The development of higher psychological processes* 5291 (1987), 157.
[57] Haiyi Zhu, Steven P Dow, Robert E Kraut, and Aniket Kittur. 2014. Reviewing versus doing: Learning and performance in crowd assessment. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*. 1445–1455.