Crowdlearning: A Framework for Collaborative and Personalized Learning

Thilak R. Balasubramanian
University of New Mexico
Albuquerque, NM, USA
trajbalasubramanian@unm.edu

Trilce Estrada
University of New Mexico
Albuquerque, NM, USA
trilce@unm.edu.edu

***Abstract*—This work presents a new technological learning paradigm that we call crowdlearning, where students are upgraded from mere passive content consumers to primary content creators and curators. Crowdlearning is all about allowing learners to share their particular vision of the world, understanding how different users learn, and putting in place mechanisms to safely, scalable, and accurately share and consume learning materials online. By providing an interactive learning framework, we expect that attention of underrepresented groups in STEM, particularly women, will be drawn towards computer science and other technology-related fields.**

In this paper we introduce all the necessary components of a crowdlearning framework, we provide a middleware implementation along with a proof of concept application. Finally, we discuss possible implications, limitations, and prospects of this approach, and how it could be shaped into a mainstream collaborative way of enhancing cyber learning. If successful in the long run, this paradigm has the potential of advancing technology aided education tools while promoting teaching, training, and collaborative learning.

I. Introduction

Online and remote education has become a priority for higher education institutions. Technology-aided educations aim is to attract and retain a broader student base. However, current remote learning technologies, such as MOOCs [1], [2], [3], [4] quickly lose student engagement due to their, often, rigid approach. Another limiting aspect of current approaches is that they rely on carefully curated content creation. Usually educators spend countless hours creating content material (e.g., videos, lectures, tests) that is then presented to students. While sometimes this process includes some form of interaction (through forums, online comments, or even some gaming aspect), students are systematically taken out of the creation cycle. This work presents a new technological learning framework that we call crowdlearning, where students are upgraded from mere passive content consumers to primary content creators. A key feature of this paradigm is its adaptability, since it can be potentially coupled to a variety of learning environments. Including traditional classroom settings: as a supplementary way of engaging student participation, to MOOCS and other on line learning modes: as an alternative way of spreading and generating content.

Through numerous examples (e.g., FoldIt [5], Exscitech [6], and other applications based on games [7]), it has been shown that involving the community in solving hard-to-model, scientific challenges provides exciting new results to the scientists and broadens the participants interest in science. For this research, we build upon these concepts; and similarly to crowdsourcing or crowdfunding, crowdlearning relies on the driving force of the crowd to achieve its goals. More specifically, *the crowd* comprises all the participants in the learning cycle, from creators and consumers of material to evaluators of other users. Then, students, educators and the public in general are part of the crowd and interact with each other in different capacities. Our paradigm seeks to reshape the traditional vertical hierarchy of knowledge transfer, where educators produce material and students consume it. Instead, both educators and students hold a horizontal relationship, where expertise and trustworthiness is conceded by the crowd.

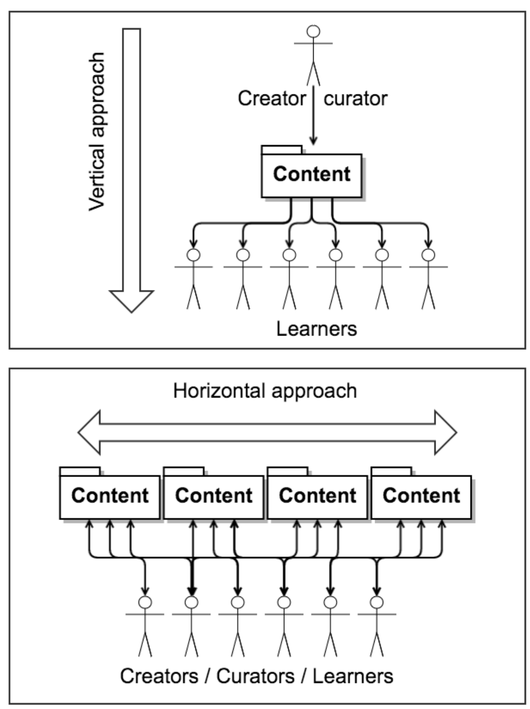

Figure 1. Vertical vs. Horizontal approaches for content dissemination

In Figure 1 we depict the difference between a top-down, vertical approach and a horizontal, crowdlearning approach for online education. In the vertical approach the responsibility of creating and curating educative material (i.e., content) lies in a reduced number of contributors; limiting effectively the scalability of this approach. On the other hand, the horizontal approach can potentially scale with the number of participants, since the creation and curation of material is a shared responsibility. Another advantage of the crowdlearning approach is that the conceptual view of educative materials does not come only from educators but also from peers. Many times the way in which a teacher explains a concept is influenced by his or her acquired experience and it is contextualized within a more complex framework than that of the students. By enabling peers to share their own ways of understanding and learning material, we can enable a more relatable way of transferring knowledge. A long term goal of this project is that by providing material with diverse perspectives and strong interactivity mechanisms, the paradigm will be appealing to underrepresented groups in STEM fields.

To make possible this paradigm, we developed the cyberinfrastructure that will be used to deploy a variety of crowdlearning applications, as well as a proof of concept multi-modal application (i.e., an application whose material ranges from images, videos, text) for geography and history learning. Note that, similarly to BOINC [8] and the Mechanical Turk [9], both examples in a different scope, the main focus of this paradigm is not developing the applications themselves, but to provide the cyberinfrastructure to make them possible and easily accessible to anyone with just basic programming skills. An important limitation of the crowdlearning approach is that the quality and trustworthiness of the material is not guaranteed. Since any participant can create content, this approach is vulnerable to malicious attackers and inaccuracies due to lack of experience. We address this drawback in one of our contributions. More generally, the contributions of this work are as follows:

- A cyberinfrastructure framework that enables easy deployment of crowdlearning applications.
- A graph theory-based approach to rate and spread user trustworthiness in a bipartite model.
- A proof of concept multi-modal crowd learning application to demonstrate some of the crowdlearning paradigm's functionality.

The rest of this paper is structured as follows: Section II presents alternate approaches for online learning and crowd sourcing. Section III describes the taxonomy of a crowdlearning application. Section IV presents a proof of concept application on geography learning. Section V discusses limitations and possible research avenues for crowdlearning. Finally VI presents conclusions and future directions of this work.

II. Related Work

Massive Online Open Courseware (MOOCs) are reshaping educational technology by democratizing access to high quality courses offered by very prestigious universities. MOOCs have evolved from being just video repositories to more mature adaptive learning models. However, since their conception, MOOCs have suffered from low retention [4]. Nowadays, MOOCs structure is diverse and learners have the option to chose from multiple categories. Retention and learners effective engagement greatly depend on different motivating factors[2], but it has been shown that more interactive MOOCs that provide environments that mix immersion and quick feedback cycles tend to be the most successful [1].

Additionally to interactivity, MOOCs have been trying to include also personalization. Work from Tovar et al.[10], proposes that MOOCs developers provide a customizable framework to be used by educative institutions to include curriculum and learning directions depending on the learners profile. Also, work from Oliveira et al., [3] suggests to look into users skills, such as the ability to self-regulate and to work in a collaborative way, and not just to replicate the content already taught in the traditional classroom, through videos or other passive learning activities. In particular, they propose conceptual strategies for educators to design MOOCs using techniques from the Flipped Classroom Teaching Model.

Our proposed framework's main difference is that educational material won't be created exclusively by the course creators/curator, but it will be directly contributed by the learners. Also, our recommendation strategies are not rigid or predefined based on users profile but they borrow concepts from collaborative filtering recommender systems. Our paradigm does not seek to replace MOOCs but it provides a scalable way of extending and enriching them, not only on the amount of material that can be produced, but also on the capability to engage and promote learners' interaction with each other.

Moreover, this research focuses on platforms that enable crowdsourcing and in its integration with learning and education [11], [12], [13], [14]. The main referent for commercial crowdsourcing is Amazon's Mechanical Turk [9], where extensive research has been done to create speech and language data [15], [16], [17] or annotate images [18] and medical entities [19]. The main drawback of commercial crowdsourcing is that these systems are plagued by uncommitted users that can potentially hinder the scientific significance of collected results. In contrast, scientific and educative crowdsourcing is likely to attract more reliable users [20].

Applications of scientific crowdsourcing include folding of proteins [5], malaria research [7], and design of experiments for docking of proteins [6]. Even though these systems have been successful, they do not expand beyond their specific domains. In particular, this work was partially inspired by [6] and [20], which presents ExSciTecH. This system uses gameplay and VC to leverage the unused power of thousands of computers to engage diverse communities in aiding scientific discovery. In ExSciTecH gameplay scenarios are built side by side to the volunteer computing (VC) mechanisms. Traditionally, scientists execute their simulations in VC projects by generating jobs, which execute in volunteers' computers when the computers are idle. The same happens in Docking@Home [21] where the scientist designs jobs by generating new input files including a protein and a ligand. The protein-ligand complexes are sent to the D@H server where they are assigned to a volunteer's client which executes them and returns results (docked ligand conformations and their energy) to the server for the scientist's analysis. ExSciTecH plugs into the BOINC framework extending it with two main gaming components: a learning component that includes a suite of games for training users on relevant biochemical concepts, and an engaging component that includes a suite of games to engage volunteers in drug design and scientific discovery. ExchiTECH shows that by enabling active participation of people in cognitive tasks beyond the passive cycle donation in paradigms like volunteer computing (VC), it is possible to use computing and mind power together in order to accelerate scientific discovery. The main difference with this work is that we won't deal with specific applications, other than a proof of concept, but will provide the general cyberinfrastructure needed to deploy them.

In previous work [20] about benchmarking gender differences in a distributed system like Volunteer Computing, in a project with over 25,000 participants, conclusions indicate a more committed response of women and ethnic minorities when the scientific component of the project is clearly disseminated. We also show that women and minorities express a stronger inclination towards STEM fields *after* participating in one of these projects. At the same time, by making VC more attractive, it is possible to influence females and ethnic minorities to consider careers in STEM fields [20]. Potential benefits to scientific dissemination, new ways to present, create, and consume knowledge, and the opportunity to engage a broader population in STEM fields motivate the work in this paper.

III. Taxonomy of a Crowdlearning Application

Before explaining what is a crowdlearning application, let us introduce the middleware's architecture (see Figure 2). Our framework consists of two main components: an analytics section centered around a data repository, and a sandbox for application subscription.

The analytics section of our middleware collects and stores user's data to perform basic actions such as authentication, along with more sophisticated analytics such

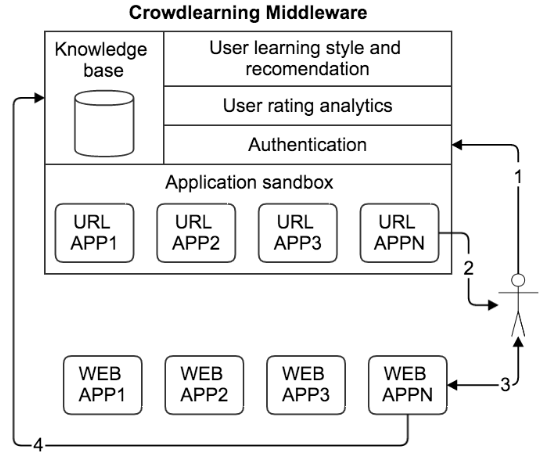

Figure 2. Crowdlearning middleware's architecture and basic user interaction. (1) The user gets authenticated and selects a particular crowdlearning application, (2) the user is assigned with a learning task and gets credentials for its execution, (3) the user interacts back and forth with the application until the task is fully executed, (4) user's interactions and metrics are saved into the middleward for further analysis

as personalized content recommendation. The information stored in the repository includes: authentication credentials, profile information, interaction patterns with the system, user's contributions, and content ratings.

The section of application subscription is a sandbox where application developers can associate their crowdlearning applications with our platform and get all the authentication, analytics, and user-base benefits of our framework. Additionally we provide an API and templates for easy deployment of new applications.

Thus, in this context, a crowd learning application is a collection of learning tasks that are related to one subject matter and may include one or all of the following phases: content creation, content evaluation, content consumption, and adaptive content release. The application can be live for a period of time or indefinitely. The time frame of the application depends on its specific goals. Applications can be created autonomously and just interface with our cyber infrastructure to get the quality control and analytics that our framework is designed to provide.

A learning task is a set of cognitive actions to be performed in tandem by a learner. A task could be (1) a tutorial, reading facts, looking at images and watching videos related to the subject matter over a period of time, (2) a quiz or test, answering to multimodal questions (e.g., questions that are presented in a variety for multimedia formats), (3) a content evaluation, rating a particular piece of content according to a variety for factors, for example rating how effective is a video or a text to explain a particular concept, (4) a didactic game, geared towards showing the use of specific concepts or to perform particular high-skilled tasks.

A user is anyone with a login account and who interacts with the application in any capacity. Any crowd learner can interact with the system in two ways: as a content creator and as a content consumer. Content creators build educative material. This creation is easy and straightforward and it is not restricted to educators; any user can create and publish learning tasks. A content consumer pulls and executes the learning tasks. That is, a content consumer would be able to perform learning activities that were published by the content creators. Once they perform the activity, content consumers can rate, evaluate, and provide feedback associated to such tasks.

A crowd learning application is divided into 5 phases:

- Application deployment: a new application is deployed to be used online. The application must have a defined set of tasks, specific goals, and a timeline if appropriate.
- User registration and authentication: users should register into the system and have valid credentials. Users are associated to a profile and a reputation. The profile and reputation are used by our crowd learning platform to assign tasks to the user. More critical tasks are assigned to users with higher reputation.
- Task execution: involves the execution of a particular learning task. Within a task execution the user consumes, and optionally creates or evaluates content.
  - Content consumption: users perform and complete learning tasks. Multimodal feedback is collected and reported to the application in order to improve it at different levels, including: content quality, information flow, multimodal aids, etc.
  - Content creation: any user can create multimodal content pieces and associate them to a particular cognitive goal in an application. More specifically, users can illustrate or define a concept or scientific process using any form of multimedia material and publish their own interpretation within a particular crowd learning application.
  - Content evaluation: once content has been created, a user can rate it based on how appropriate it is for the intended audience, and how effective the material is to convey a particular cognitive piece of information to the rating user specifically. Content evaluation goals are twofold: first, to ensure the credibility and high quality of the published material; second, to gather specific users' preferences and learning styles, that can be used later to provide personalized content recommendations.
- Adaptive content release: consists on analyzing task completion and their effectiveness in engaging the crowd to consistently contribute to improve content. Also, specific information from user's interaction with the platform is used to personalize content recommendations.

A. Application deployment

Our framework was designed to enable maximum flexibility between crowd learning applications. We provide an API to keep track of the authentication and analytics required by the applications, but the details and particular implementation can be customized to the specific needs of each project.

Representational State Transfer (REST) has emerged in the last few years as a predominant Web service design model. REST applications are easy to integrate in the web as they use HTTP methods explicitly, all resources are identified by URIs, which provides a simple way to deal with the evolution of a system, and they can achieve security through both the transport layer (SSL) and a variety of message-level mechanisms. Deploying a crowd learning application into our framework is easy and straightforward. Developers need to, optionally, host their Web application and activate two REST services. The first REST service enables authentication and profile association of users. The second REST service retrieves user's performance, which is later used for analytics and personalized recommendations.

In general, a crowd learning application must have a defined set of tasks, specific goals, and a timeline if appropriate. Each task contains learning elements (e.g., answering questions, assimilating some material, rating content) and is associated to cognitive goals (e.g., being able to identify a concept, being able to relate two or more concepts, being able to generalize a concept from a set of particular examples). Our framework provides templates and placeholders to easily deploy and extend all these different types of learning tasks into a crowd learning application. Also, we provide metrics associated to each cognitive goal and how they relate to the different learning tasks. We store and report metrics so that they can be analyzed and used to improve the learning tasks, to measure their effectiveness and also to characterize and personalize content for the users. Our metrics are classified as temporal, performance, and engagement.

- Temporal metrics include timed responses (i.e., how long does it takes to answer a question), time spent on a particular material (e.g., time watching a video or reading a text) and they provide us with a progression of the learners in terms the expertise they gain and the interest they show towards some material.
- Performance metrics include accuracy of their responses and material completion rates; these metrics are straightforward and measure proficiency of the learner over time.
- Engagement metrics include number of user contributions watched, rated, and suggested; these metrics measure how invested is the user in the application.

The application can be up for a period of time or indefinitely. The time frame of the application depends on its specific goals. Applications can be created autonomously and just interface with our cyber infrastructure to get the broader user base, quality control, and analytics that our framework is designed to provide, or can be created using our templates and dedicated functionality.

B. User registration and authentication

User authentication is a very important aspect of our framework for two reasons: (1) it is important to be able to differentiate every user from the rest, so that they are linked to their specific profile, thus anonymous access is not allowed; (2) it is crucial to try to stop malicious users and avoid boot attacks. Therefore, we use a two level authentication process. The first level is just a traditional credential matching, where users input a login and password. Once an user has been authenticated past this first level, it is possible to retrieve their profile, scores, and related metrics. Using this personalized report, our middleware presents the user with a list of tasks that he or she can perform. The list of tasks is defined by the user's proficiency, past behavior, and preferences. The user can select from this list which task to perform.

Once the user has selected a task, the second level authentication is activated. Our framework generates a random 6-digit one-time password (OTP) to grant the execution of that specific task by the specific user. An OTP is a password that is valid for just one login session; the advantage of OTP passwords is their limited temporality. Contrary to static passwords OTPs are not vulnerable to replay attacks and can effectively contain boot attacks.

C. Content consumption and creation

Once the user has been authenticated and a task has been assigned to him or her, the user can then proceed to execute it. In the event that the user would like to recommend better material, he or she can upload content. Finally, the user can rate how effective was the multimodal content to perform the specific task.

Content consumption refers to actually performing the specific learning task. This can include answering a test, going through a tutorial, or reviewing random material associated to a particular topic. Many of these tasks would be timed or would keep track of completion rates (as mentioned before in the performance metrics).

Through the execution of a task, the user will receive multimodal feedback. For example, if the user is going through a tutorial, he or she would be presented first with a text, if the concept is clear (i.e., if the user rates highly), then they can continue with the next section. If the text is not clear, the user will be presented with a set of images (for example, obtained from Instagram using tags from the text) or videos (obtained from YouTube in a similar way). The user can rate the effectiveness of the feedback and optionally provide comments. In addition to the ”subjective” rate, we can measure objectively the user's performance by keeping track of their future performance on the topic.

If a particular task execution fails because of technical reasons (e.g., the user gets disconnected, the application is out of reach) then the task goes back to the pool of executable tasks and the user can try to execute it again at a later time. Partial progress can be recorded in the application side and in this case work recovery could be implemented. However, this feature is not enforced and it is up to the application to provide recovery or not. In general many applications would not allow partial recovery and they will be notified of material that was accessed only partially. Reporting is useful to avoid unfair advantages on timed quizzes, for example.

Content creation can occur in different phases. The most straightforward way is when application creators build or link content into their specific cognitive tasks. However, the most interesting way, and what makes this paradigm scalable and horizontal, is when learners provide feedback and contribute content to tasks. The amount of material associated with a task can potentially grow proportionally to the number of users. Every user can provide a different view to a specific problem. The main advantage of this approach is that it promotes diversity of learning styles. The educative material that works best for a set of users may be very different for another set. Thus, we called this paradigm Crowd Learning because it takes advantage of individual contributions to better serve the crowd of learners as a whole.

D. Content evaluation

Content evaluation goals are twofold: first, to ensure the credibility and high quality of the published material, second, to gather specific users preferences and learning styles, to be used later to provide personalized recommendations. Contrary to MOOCs or other paradigms where content is carefully made and curated, crowdlearning has an issue with trustworthiness and accuracy. Users can create content and publish it freely. However, this freedom can be potentially abused and users could create offensive or unrelated material. Also, because of the different levels of proficiency, users may create incomplete or inaccurate content. Thus, an evaluation mechanism is needed to ensure the quality and integrity of the learning material.

To perform content evaluations, users are associated with a profile and a reputation score. The reputation score is earned over time and interaction with the platform. The user hierarchy is defined initially by their association with crowdlearning applications. When a new application is created it could have a set of super users who are preemptively approved to curate material and they would have a large trustworthiness score for that application. A new user, that does not have the rank of superuser, would start with a reputation score of zero; meaning that the user is not trusted to begin with. However, as its contributions are approved by

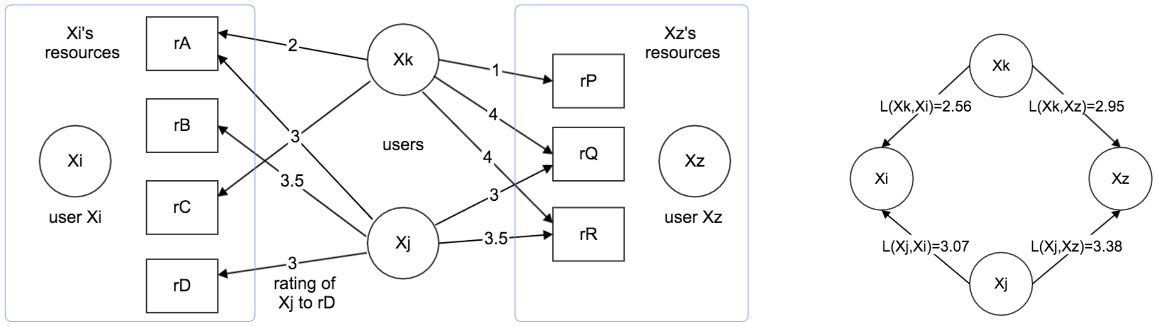

(a) Representation of a bipartite graph for users, resources, and ratings

(b) Calculation of node links in a user's graph

Figure 3. Modeling trustworthiness as a directed graph

other users with higher reputation, his or her reputation score improves too.

We model the trustworthiness hierarchy as a graph problem, where ratings from one user to another user's contributions are added as weighted links between nodes representing users. Figure 3(a) represents users as labeled circles, resources are represented as rectangles and they can be videos, images, and text contributed by a particular user. Links denote ratings of users to resources. Multiple users can rate multiple resources and a resource can be rated by multiple users but it belongs to one single user. Figure 3(b) shows the corresponding consolidated graph where only users are linked to each other. The strength of the link is calculated based on the ratings of a particular user to the resources of another one.

Consider users x_i and x_j from Figure 3 and assume that r_j is x_j's rating vector. To calculate the weight of a link between user x_i and x_j we need to determine (1) what is the standard used by x_j to rate contributions (given by $\sum_n (r_j^k - \mu(r_j))/\sigma(r_j)$), (2) what are the scores of the contributions of x_i that x_j has rated, and (3) how many total rating has x_j made (given by $|r_j|$). Thus the weight of a link from x_j to x_i is:

$$L(x_j, x_i) = \mu(r_j) + \frac{\sum_k \left(\frac{r_j^k - \mu(r_j)}{\sigma(r_j)} \right) \delta(k, x_i)}{|r_j|} \quad (1)$$

Where r_j is the vector of size $|r_j|$ containing all the ratings made by x_j, $\mu(r_j)$ and $\sigma(r_j)$ are its mean and standard deviation respectively. r_j^k is the specific rating given by x_j to resource k, and $\delta(k, x_i)$ is a delta function that returns 1 if resource k belongs to x_i and returns 0 otherwise. With this definition of a link weight we ensure that every pair of users are linked in a fair manner. That is, in a way that takes into account the popularity of a user's resources and the different value scales used by the different users to rate resources.

The overall user reputation score is then calculated using PageRank [22], where the reputation scores are propagated in proportion to how contributions from users have been rated by the most trusted superusers. More specifically, if a node x_j has a total score of $score(x_j)$ and its outgoing links are weighted as shown in Equation 1 as $Link(x_j, x_n)$, then the score that is propagated from node x_j to x_i is given by:

$$score(x_j, x_i) = \frac{L(x_j, x_i)}{\sum_n L(x_j, x_n)} * score(x_j) \quad (2)$$

Figure 4 shows a fragment of a trustworthiness network centered around a superuser with a total score of 100. The score of the super user is calculated as the sum of the scores transferred by rating her contributions. For simplicity purposes we show only a few nodes and just the ego network of the superuser. At the same time, the superuser transfers her score proportionally to nodes to which she has rated their contributions.

Note that this hierarchy is only designed to validate user behavior and not to restrict it in any way. Untrusted users can still create content and eventually improve their scores, however their contributions will be more strictly scrutinized than contributions from already trusted users. This user hierarchy and rating system, provides crowdlearning with its own regulation mechanism. Still, in order to prevent that students learn inaccurate information, strict quality control of content will be ensured by redundantly validating content through the user hierarchy, and releasing information to more inexpert users only when the content has been throughly approved by more experienced users.

E. Adaptive content release

The second motive to performing content evaluation is to uncover specific user's learning preferences and take

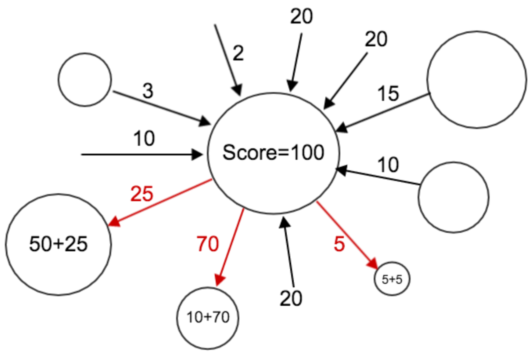

Figure 4. Score transfer through PageRank in an egocentric user network

advantage of such preferences and particular style to improve the mechanism that assigns tasks among users.

We explore the learning styles theory [23], [24] and propose a method to suggest more relevant and interesting material to users. In particular we borrow techniques from the Big Data community to filter information based on distinctive patterns between multiple data sources. We use collaborative filtering [25] to find what type of content is rated highly by sets of users and then use this association to recommend material that is also likely that they will rate positively. The analytics aspect of our framework is in charge of aggregating user's behavior and relating it to specific resources and resource types. The deployment and evaluation of this mechanism is work in progress.

IV. Proof of Concept: A Multimodal Crowd Learning Application on Geography

To show some of the features of the crowdlearning approach and the easy mechanisms in place to couple applications with our middleware, we developed a proof of concept application. This application, for learning national geography at a K-12 level, contains two learning tasks: a multi-modal tutorial describing some major geographic landmarks, and a quiz that evaluates the user's understanding of some of those landmarks. Figure 5 shows a fragment of this application deployed in a UNM server.

The application was developed in Java with a basic functionality front end on JavaScript and user interface based on BootsrapCSS. The back end has a MySQL database to keep track of the material and user contributions and uses JDBC to interface with the front-end. The application was built using an MVC pattern with business objects (BOs) and data access objects (DAOs).

All material can be directly saved into the database by the application developpers. At the same time, all the material can be pulled as if it was a user contribution just using keywords and hashtags. Text can be retrieved from

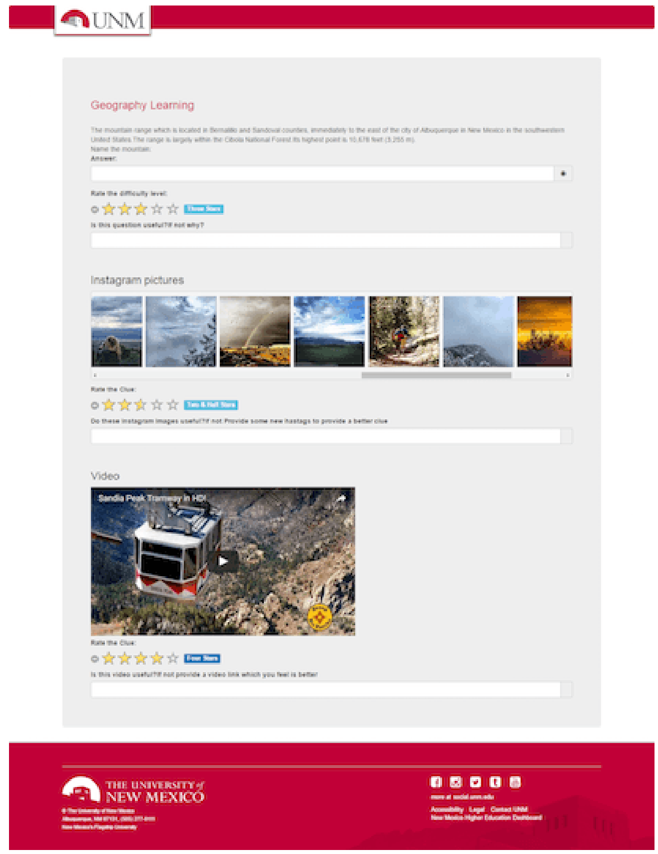

Figure 5. Proof-of-concept crowd learning application in geography

Wikipedia [26], images from Instagram [27] and videos from YouTube [28], just to cite some examples.

To keep track of some user's metrics, particularly the timing metrics, we used the *onLoad* function of JavaServer Pages (JSP), which keeps track of the time from the load of the page up until the *Next* button is pressed. For videos, since it is important to track whether a video was watched in its totality (e.g., to measure engagement and effectiveness of its content) we embedded the videos using the *iframe* tag in HTML5 and use its *watchTime* report.

The application itself is a template for tutorials and quizzes displaying multimodal content that can be rated and augmented with comments or suggestions. We developed this application just as a proof of concept and as a way to illustrate the use of our templates for the creation of similar applications. In this context, the subject matter can be easily changed, but the functionality can be reused for any other theme.

V. Beyond a Proof of Concept: What Crowdlearning Could Bring to the Table

A. *Enriching traditional and online learning environments*

The ultimate goal of this paradigm is to be used as an aid or complement in a variety of learning environments. We envision that traditional classroom settings would be able to use it for student's homework or evaluations, and that MOOCS and online tutoring systems will benefit from it as a way to diversify and generate additional content.

The produced data obtained from participant's interactions, will be annotated in order to be stratified, while still preserving the privacy of the participants. The data will be made available to the public and researchers to enable socio-educative studies. At the end, this project will become a Big Data repository capable of producing insights on learning modes and preferences as well on material effectiveness and best practices of how to couple this approach into other educative environments.

B. Providing a platform for crowdsourcing highly-skilled cognitive tasks

Crowdlearning is not only a platform for hosting high school quizzes, tutorials, and games. Its capabilities can go beyond sharing and disseminating knowledge and can become a trusted platform for donating highly skilled mental abilities for solving or contributing to scientific problems. We envision that crowdlearning can eventually provide scientific applications with an infrastructure to deploy active learning through crowdsourcing. While the scientific value of crowd contributions needs to be thoroughly tested and validated, its community-building and engagement properties are likely to show positive outcomes almost immediately. The basic idea is to involve citizens into scientific projects to promote their interest in science while learning some basic concepts. When a person takes ownership of a contribution and feels part of a cause, they become more involved in the foundations of said cause [29], [30]. In this particular case those foundations include mathematics, chemistry, physics, medicine, and computer science. Our goal is to indirectly increase awareness of these fields through what appears to be mere games and pastimes. If successful in the long run, this paradigm has the potential of advancing technology aided education tools while promoting teaching, training, and collaborative learning.

C. Attracting underrepresented groups to STEM fields

The involved approach to education and hihly diverse content of corwdlearning has the additional advantage of potentially attracting women and minorities to STEM through the scientific component of this framework. Although there is no general consensus on the learning styles' theory as a whole, it has been proposed that women show a marked inclination to be, what is called kinesthetic learners [23]. A kinesthetic learning style involves seeing, hearing, doing and experiencing things [24]. From the science and engineering perspective, this learning style requires putting abstract concepts into a real-life context. For example, a kinesthetic learner benefits from play-acting or puppet shows at a small age and from group interaction and immersive experiences when they are older. By providing an interactive learning framework, we expect that attention of women will be drawn towards computer science and other technology-related fields.

VI. Future Work

Work in progress from the development of this framework includes the successful integration of analytics and collaborative filtering recommendation into our framework.

More importantly, logistic future work includes reaching out to educative institutions with a large user base and educative content needs. Our plans include establishing a partnership with middle and high education institutions and to deliver proof of concept applications to spark their initial interest. Our first target will be simple physics and chemistry crowdsourcing tasks that we plan to release as games and educational exercises for students. The application will revolve around basic concepts and will provide an interactive platform for science teachers in high schools to convey concepts more effectively and to quantitatively assess student participation. As a side effect, we expect that by interacting with these problems, young students will be drawn to STEM fields.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation for the grant entitled CAREER: Enabling Distributed and In-Situ Analysis for Multidimensional Structured Data (NSF ACI-1453430).

References

[1] M. Freire, I. Martnez-Ortiz, P. Moreno-Ger, and B. Fernndez-Manjn. Requirements for educational games in moocs. In *2015 IEEE Global Engineering Education Conference (EDUCON)*, pages 993–997, 2015.

[2] D. G. Sooryanarayan and D. Gupta. Impact of learner motivation on mooc preferences: Transfer vs. made moocs. In *2015 International Conference on Advances in Computing, Communications and Informatics (ICACCI)*, pages 929–934, 2015.

[3] A. G. de Oliveira Fassbinder, M. Fassbinder, and E. F. Barbosa. From flipped classroom theory to the personalized design of learning experiences in moocs. In *Frontiers in Education Conference (FIE), 2015. 32614 2015. IEEE*, pages 1–8, 2015.

[4] J. Xiao, B. Jiang, Z. Xu, and M. Wang. The usability research of learning resource design for moocs. In *Teaching, Assessment and Learning (TALE), 2014 International Conference on*, pages 277–282, 2014.

[5] A. L. Beberg, G. Jayachandran, S. Khaliq, and V. S. Pande. Folding@home: Lessons from eight years of volunteer distributed computing. In *In Proceedings of the 8h IEEE International Workshop on High Performance Computational Biology*, pages 3–16, 2009.

[6] M. Matheny, S. Schlachter, L. M. Crouse, E. T. Kimmel, T. Estrada, M. Schumann, R. Armen, G. Zoppetti, and M. Taufer. Exscitech: Expanding volunteer computing to explore science, technology, and health. *2012 IEEE 8th International Conference on E-Science*, pages 1–8, 2012.

[7] John A. F. Reliable enumeration of malaria parasites in thick blood films using digital image analysis. *Malaria Journal*, 8, 2009.

[8] D. P. Anderson and K. Reed. Celebrating Diversity in Volunteer Computing. In *Proc. of the Hawaii International Conference on System Sciences (HICSS)*, 2009.

[9] Amazon mechanical turk. https://www.mturk.com/mturk.

[10] A. Alzaghoul and E. Tovar. A proposed framework for an adaptive learning of massive open online courses (moocs). In *2016 13th International Conference on Remote Engineering and Virtual Instrumentation (REV)*, pages 127–132, 2016.

[11] V. Ambati. *Active Learning and Crowdsourcing for Machine Translation in Low Resource Scenarios*. PhD thesis, 2012.

[12] Y. Cheng, Z. Chen, L. Liu, J. Wang, A. Agrawal, and A. Choudhary. Feedback-driven multiclass active learning for data streams. In *Proceedings of the 22nd ACM international conference on Conference on information and knowledge management*, pages 1311–1320. ACM, 2013.

[13] M. Georgescu, D. D. Pham, C. S. Firan, U. Gadiraju, and W. Nejdl. When in doubt ask the crowd: Employing crowdsourcing for active learning. In *Proceedings of the 4th International Conference on Web Intelligence, Mining and Semantics (WIMS14)*, pages 12:1–12:12. ACM, 2014.

[14] J. Pujara, B. London, and L. Getoor. Reducing label cost by combining feature labels and crowdsourcing. In *ICML Workshop on Combining Learning Strategies to Reduce Label Cost*, 2011.

[15] Q. Gao and S. Vogel. Consensus versus expertise: A case study of word alignment with mechanical turk. In *Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon's Mechanical Turk*, pages 30–34. Association for Computational Linguistics, 2010.

[16] J. Gordon, Van D., Benjamin, S., and Lenhart K. Evaluation of commonsense knowledge with mechanical turk. In *Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon's Mechanical Turk*, pages 159–162. Association for Computational Linguistics, 2010.

[17] C. Akkaya, A. Conrad, J. Wiebe, and R. Mihalcea. Amazon mechanical turk for subjectivity word sense disambiguation. In *Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon's Mechanical Turk*, pages 195–203. Association for Computational Linguistics, 2010.

[18] C. Rashtchian, P. Young, M. Hodosh, and J. Hockenmaier. Collecting image annotations using amazon's mechanical turk. In *Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon's Mechanical Turk*, pages 139–147. Association for Computational Linguistics, 2010.

[19] M. Yetisgen-Yildiz, I. Solti, F. Xia, and S. R. Halgrim. Preliminary experience with amazon's mechanical turk for annotating medical named entities. In *Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon's Mechanical Turk*, pages 180–183. Association for Computational Linguistics, 2010.

[20] Trilce Estrada, Kathleen L. Pusecker, Manuel R. Torres, Joanne Cohoon, and Michela Taufer. Benchmarking gender differences in volunteer computing projects. In *Proceedings of the 2013 IEEE 9th International Conference on e-Science*, pages 342–349. IEEE Computer Society, 2013.

[21] M. Taufer, R.S. Armen, J. Chen, P.J. Teller, and C.L. Brooks III. Computational multi-scale modeling in protein-ligand docking. *IEEE Engineering in Medicine and Biology Magazine*, 28:58–69, 2009.

[22] L. Page, S. Brin, R. Motwani, and T. Winograd. The pagerank citation ranking: Bringing order to the web. *Stanford Digital Libraries Working Paper*, 1998.

[23] J. I. A. Alumran. Learning styles in relation to gender, field of study, and academic achievement for bahraini university students. *Individual Differences Research*, 4:303–316, 2008.

[24] McDaniel M. Rohrer D. Bjork R. Pashler, H. Learning styles: Concepts and evidence. *Individual Differences Research*, 9:105119, 2009.

[25] Michael D. Ekstrand, John T. Riedl, and Joseph A. Konstan. Collaborative filtering recommender systems. *Found. Trends Hum.-Comput. Interact.*, 4(2):81–173, February 2011.

[26] Wikipedia, the free encyclopedia. https://www.wikipedia.org/.

[27] Instagram. https://www.instagram.com.

[28] Youtube. https://www.youtube.com/.

[29] A. S. Brown and M. Oxman. Learning through participation: Effects of involvement and anticipation of involvement. *The American Journal of Psychology*, 91(3):461–472, 1978.

[30] B. Lundell, A. Persson, and B. Lings. Learning through practical involvement in the oss ecosystem: Experiences from a masters assignment. 234:289–294, 2007.