

Flexible Learning with Semantic Visual Exploration and Sequence-Based Recommendation of MOOC Videos

Jian Zhao Chidansh Bhatt Matthew Cooper David A. Shamma
FX Palo Alto Laboratory (FXPAL), Palo Alto, CA 94304 USA
{zhao, bhatt, cooper, ayman}@fxpal.com

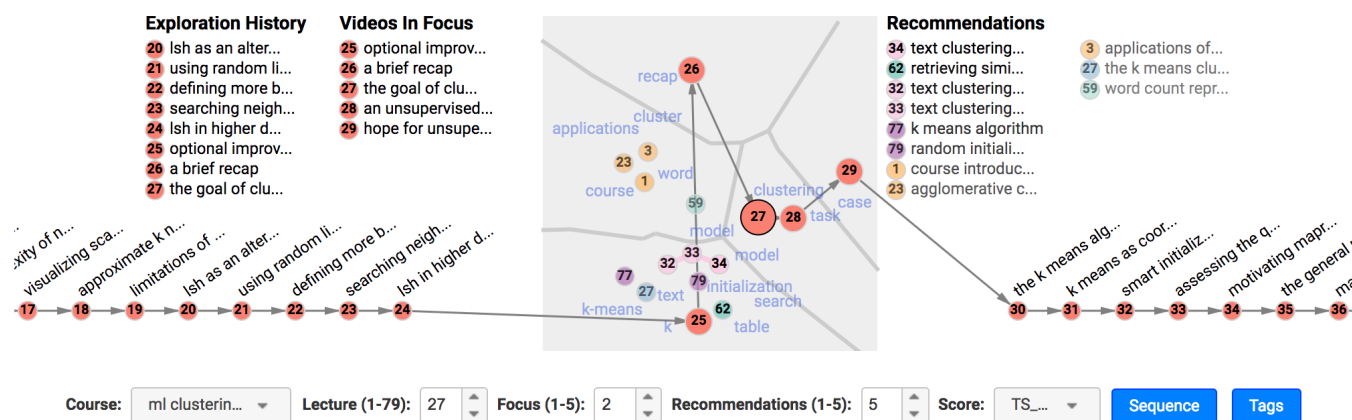


Figure 1. A learner is using MOOCex to explore recommendations of lecture video 27: *The goal of clustering in course Machine Learning: Clustering and Retrieval*. Videos are color-coded by course and labelled with lecture sequence numbers. Videos from the current course are connected with arrows. The recommendations and the local neighborhood of the current video are presented in a 2D space based on their relationships. The space is further split into multiple regions according to coherence in semantics, and is augmented with keywords extracted from the video transcripts.

ABSTRACT

Massive Open Online Course (MOOC) platforms have scaled online education to unprecedented enrollments, but remain limited by their rigid, predetermined curricula. To overcome this limitation, this paper contributes a visual recommender system called MOOCex. The system recommends lecture videos across different courses by considering both video contents and sequential inter-topic relationships mined from course syllabi; and more importantly, it allows for interactive visual exploration of the semantic space of recommendations within a learner’s current context. When compared to traditional methods (e.g., content-based recommendation and ranked list representations), MOOCex suggests videos from more diverse perspectives and helps learners make better video playback decisions. Further, feedback from MOOC learners and instructors indicates that the system enhances both learning and teaching effectiveness.

ACM Classification Keywords

H.5.2. Information Interfaces and Presentation (e.g. HCI): User Interfaces

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.
CHI 2018, April 21–26, 2018, Montreal, QC, Canada.
© 2018 Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-5620-6/18/04...\$15.00.
DOI: https://doi.org/10.1145/3173574.3173903

Author Keywords

MOOC Videos, Visualization, Recommender System.

INTRODUCTION

Modern online education platforms, such as Coursera, edX, and Udacity, have become increasingly popular in recent years. These platforms allow for teaching at a distance and at scale by presenting educational materials as Massive Open Online Courses (MOOC). A course usually consists of a number of short videos, each targeting a specific concept. To achieve certain learning objectives, instructors commonly order the videos according to a syllabus which may also group videos hierarchically into sections.

However, the syllabus remains a *one size fits all* approach with a predefined curriculum, which contributes as a critical factor in courses’ low retention rates [38, 54]. Further, studies show that professionals, who comprise an increasing portion of MOOC learners, aim to advance their career growth (rather than obtaining a certification), and are less likely to follow the syllabus [54]. It is critical to offer learners more flexible access to a broader range of content and perspectives (e.g., from multiple courses) [9, 27, 46].

Some platforms such as KhanAcademy [25] and booc.io [44] provide an interactive knowledge (concept) map or visualization that allows for more personalized learning behaviors. However, concept maps are not well suited for sequential flow [37] and creating a dependency of concepts requires

manual effort; this is neither scalable nor adaptive. We tackle this problem by offering recommendations in appropriate visualization to enable learners to access content with multiple perspectives from related courses available on different platforms. Learners can interactively navigate recommended videos in the semantic space of their learning context, meanwhile not becoming *lost in hyperspace* [16].

In this paper, we present MOOCex, a system that facilitates effective browsing of relevant content based on interactive visualization and on-the-fly video recommendation (Figure 1). Built upon advanced data mining techniques, MOOCex offers diverse perspectives on related concepts through recommendation of video sub-sequences that considers both videos' topics and sequential inter-topic relationships between videos. Unlike conventional user interfaces for showing recommendations, i.e., a ranked list or a set of ranked lists, MOOCex supports semantic visualization of recommended videos in users' current learning context, by projecting videos onto a 2D space annotated with topical regions and key phrases. This provides additional dimensions for learners to effectively explore related content and select what to watch next wisely.

From a dataset of 4,186 MOOC videos, we conduct three experiments to evaluate MOOCex from multiple aspects, including recommendation quality measurements, interviews with MOOC instructors, and controlled studies with MOOC learners. Our goal here is to gain a holistic view of the effectiveness of MOOCex as a visual recommender system. The results of comparing MOOCex with the traditional methods (i.e., content-based recommendation and ranked lists) indicate its effectiveness and usefulness, as well as, its promising potential for applications on MOOC platforms.

RELATED WORK

In this section, we review the state of the art including video and content recommendation techniques, approaches for visualizing search, recommendation, and information spaces, as well as visual analytics systems for MOOC data.

Video and Content Recommendation

The majority of approaches to video recommendation are content-based, for example, through performing textual analysis of transcripts and visual analysis of image frames. VideoReach provides recommendation by combining three models based on textual, visual, and aural information of videos [34]. TalkMiner is an educational video search system using OCR and lexical analysis of text displayed in video frames [1]. Topic modeling, such as Latent Dirichlet Allocation (LDA) [4], has also been applied [55].

However, none of these methods has considered the sequential relationships between videos, which are available on MOOC platforms. The video sequence in a syllabus is created by instructors to achieve certain learning objectives. Recommendations ignoring this information may result in fragmented learning, thus degrading the quality of recommended content.

Most applications of sequential pattern mining techniques are distinct from video recommendation. In the field of learning analytics, Kinnebrew et al. employed sequence mining for

modeling student behavior [29], but they did not focus on content-based video processing. Morales et al. [35] facilitates information discovery via educational hypermedia linking based on sequence mining of user logs. Agrawal et al. [2] proposed a system for linking web videos to supplement electronic textbooks, and argued that topic mining alone is insufficient. More recently, Doroudi et al. indicated that topic sequences play a critical role in student performance in the context of a tutoring system [15]. In other applications, sequential information has been used in recommending music [22], online products [6], and travel itineraries [24].

The above studies demonstrate that sequential organization of topics by experts in course syllabi can provide valuable information for educational content recommendation. Our approach to recommendation combines topic modeling with sequential pattern mining to enhance recommendation quality. However, we face a bigger challenge in MOOC video recommendation: data sparsity, because videos in the syllabi normally do not overlap across different courses. This contrasts with the datasets used in previous work (e.g., music playlists [22]). To address this issue, we apply top- K sequential rules (TKS) [17] and top- K non-redundant sequential rules (TNS) [18] to analyze both global and local patterns in topic transitions exhibited in instructors' syllabi, whereas the above techniques only consider the global sequential information.

Visualization of Search, Recommendation, & Info Spaces

Our work relates to the visualization of search and recommendation results, and more generally information spaces. At a concept level, Dörk et al. introduced *information flaneur* [12] and *monadic exploration* [13] that guide interface design for information seeking.

One way of presenting the retrieved information (e.g., from search or recommendation) is based on the traditional linear form like ranked lists on Google and YouTube. For example, TileBars places a colored bar next to each list entry to convey document length and term frequency [23]. PubCloud augments a search result list with tag clouds [32]. LineUp offers rich interaction to allow users to visually manipulate a ranked list [21]. Based on list representations, ExplorationWall [31] and uRank [11] investigate methods for refining search results on-the-fly as information needs evolve.

Many techniques utilize a 2D space to present the information based on various layout methods. One approach is to place items based on their attributes (or facets). For example, PivotPath [14] and PivotSlice [53] display an academic publication corpus as a network and position each document interactively in a partition of a 2D canvas based on meta-data. InterAxis allows for manipulating a scatterplot with axes dynamically defined by users [26]. Additionally, space-filling techniques such as treemaps and 2D tiles are used for browsing searching results [8, 30]. On the other hand, WebSearchViz [36] and Topic-Relevance Map [40] employ a circular layout, where the distance to the center represents a document's relevance to a query and different sections of the circular area denote different topics. Similarly, RankSpiral employs a spiral layout to render a ranked list by maximizing information density [47].

A vast majority of 2D methods employ a more “free-formed” layout based on some similarity measure, in which the distance between items encodes their similarity. For example, Gomez-Nieto et al. proposed an energy-based method to place text snippets of search results with minimal overlap [20]. Dimension reduction techniques such as multidimensional scaling (MDS) [48] are often employed to reveal clusters of coherent content within data (e.g., IN-SPIRE [50]), and are sometimes applied in combination with topic modeling [7, 28]. Similarly, some recommender systems visualize query results in logical sets or clusters on a 2D canvas, allowing users to interactively drive recommendation strategies and incrementally explore information [5, 39, 49].

However, none of the above techniques has incorporated sequential relationships between documents, a critical aspect of MOOCs, for recommendation/search models or for visualization. An exception might be TimeCurves [3] that visualizes the sequential order of data points on top of a MDS projection, but this work does not focus on either educational materials or recommendation.

Visualization of MOOC Data

Many techniques have been proposed for visual analysis of data generated by MOOCs, such as user clickstreams and forum discussions [41]. One main topic in this area is to study learner behaviors. For example, Coffrin et al. employed state transition diagrams to visualize video switching patterns of learners. More sophisticated systems allow analysts to compare different user groups and examine detailed video navigational behaviors [10, 45]. In addition, iForum provides another perspective for understanding learners via the analysis of the content and structure of MOOC forum threads [19].

In contrast with these systems that are specifically designed for instructors or analysts, we focus on recommendation and visual exploration of MOOC videos to benefit ordinary learners. One particular work that shares similar goals with ours is booc.io [44], which allows for visually exploring concept maps of instructional materials and following personalized learning plans. However, this approach is less scalable or flexible because of the required manual creation of the concept maps beforehand. Also, they focus on video exploration within a single course whereas we recommend videos sourced across multiple MOOC platforms.

DESIGN GOALS

Our aim is to enable MOOC learners to flexibly access course content across MOOC platforms using interactive recommendation. We identified the following design goals to guide the development of our system.

G1: Facilitate decision making. Our primary goal is to allow learners to quickly and wisely choose what to watch next based on their current context. It is critical to facilitate both their understanding of unknown information spaces using semantics, and their determination of promising directions to explore further [31, 50]. As it is impossible to present the entire space (often huge and complicated), we aim to encourage curious and creative information seeking to “see the whole through its parts” [12, 13].




Figure 2. System architecture overview of MOOCex.

G2: Provide context. Different from other scenarios, learning builds upon comprehension of unit concepts in meaningful orders. Studies show that the sequence of topics is critical in affecting learners’ performance [15]. Schwab et al. advocated that presenting the dependency of concepts is essential for MOOC platforms [44]. Thus, recommender systems for educational materials should provide context for effective learning by considering both sequential relationships and content-based relevance.

G3: Offer diversity and flexibility. MOOC platforms should support diverse goals and provide flexibility for learning [54]. Studies also affirm that students perform better when provided with multiple perspectives [46]. Hence, recommended content should be diversified to offer flexibility in selecting what to view, but at the same time preserving sufficient context [16, 44].

In addition to these primary design goals, we aim to provide a compact, simple, and easily understood user interface for exploring recommendations. In general, we want users to focus on learning concepts and watching videos, rather than expending too much effort inspecting the information space.

MOOCex TECHNIQUE

We approach the above goals using recommendation combining both topical and sequential information from videos and enabling semantic visual exploration of the recommended results. In this section, we first give an overview of MOOCex and describe our experimental dataset. Next, we introduce technical details of its recommendation engine and visual interface, and finally describe a couple of use cases.

Overview and Data Collection

The MOOCex architecture consists of two main components: a recommendation engine and a visualization module (Figure 2). For recommendation, offline pre-processing of the data is conducted beforehand to derive frequent topic transition rules, and at querying time, MOOCex uses the mined information to recommend relevant video sub-sequences. The visualization serves as an interface for learners to select videos to view, as well as a tool for them to better understand the recommendations from different courses and to effectively explore the unknown huge information space.

We gathered courses available from MOOC platforms including Coursera, edX, and Udacity. For this study, we chose mainly computer science courses such as *Applied Machine Learning*, *Text Mining*, and *Data Visualization*. In total, we collected 4,186 videos from 41 courses (approximately 344.65 hours of running time). In addition to the source videos, we

collected text transcripts and course syllabi; together with the videos, it makes a total dataset footprint of 126 GB.

Recommendation Engine

The recommendation engine of MOOCex contains two stages (Figure 2). We first pre-process the videos to build the knowledge base, which contains vector representations of the videos and their sequential relationships mined from the syllabi. Then, recommendations are retrieved according to similarity in the vector space and further re-ranked with the sequential information.

Pre-Processing

In the pre-processing stage, we first develop a feature representation for each video based on the text transcripts within our dataset. We use LDA [4] globally to represent each video using its distribution over the discrete set of latent topics with dimensionality $Z = 30$ (i.e., the number of topics). We denote this as the *topic signature* of the video, $V_i = \{k : p^{(i)}(z_k) \geq 0.1\}$, which contains the indices of topics with probability for the i^{th} video of at least 0.1.

Given a consistent representation of videos, we discover sequential inter-topic relationships. Our goal is to use the currently watched video to predict likely topics users will watch next and incorporate this information in our recommendation. We employ sequential rule mining techniques to discover transition rules from one set of topics to another. We construct sequences of topics according to the order of videos in a course syllabus with section-based partitions, which generates 537 sequences from our database. We then apply the top- K non-redundant sequential patterns (TNS) [18] to detect prevalent patterns, which reflects a *global* analysis of topic sets and sequential transitions. However, some topic sets observed in our videos are relatively infrequent and largely absent from the global analysis. Thus, we employ the top- K sequential pattern mining (TKS) [17] to extract patterns with lengths between 3 and 6, for our *local* analysis of the sequences.

Recommendation

In the query stage, MOOCex provides recommendations based on the current video being watched. Content-based recommendation is first applied using standard TF-IDF retrieval based on video transcripts and cosine similarity [33, 43]. This initial ranking finds relevant videos with similar content. We then perform re-ranking based on the topic similarity score (TS), and then global (GS) and local (LS) sequential rules.

Topic similarity score (TS). The topic distributions of the query video $P^{(q)}$ and the recommended video $P^{(r)}$ are matched in terms of overlap and probability values. The TS score combines the Jaccard similarity, $J(\cdot)$, between the topic signatures, V_q and V_r , and the number of probabilities for common topics that are within a threshold t :

$$S_{TS}(P^{(q)}, P^{(r)}) = J(V_q, V_r) + \frac{1}{Z} \sum_{z=1}^Z \delta(|P^{(q)}(z) - P^{(r)}(z)| \leq t),$$

where $P^{(i)}(z)$ is the probability of latent topic z for video i and $\delta(\cdot)$ is the Dirac delta function. We set $t = 0.2$ in

our experiments. Here, the first term measures the number of common topics between two videos, and the second term captures the topic distributional similarity. This score emphasizes videos with similar concepts from the perspective of other courses. As both terms are related to topic similarity, future studies are warranted to examine their correlations.

Global sequence score (GS). We retrieve N support and confidence score values, $\{(s_i, c_i)\}$, from the mined global sequential patterns with antecedent values matching V_q and consequent values matching subsets of V_r . The GS score is

$$S_{GS}(V_q, V_r) = \frac{1}{N} \sum_{i=0}^N c_i \frac{y_i}{|V_r|} + \frac{s_i}{D_G},$$

where D_G is the total number of global sequences, and y_i is the length of the matched subset of V_r . To avoid noisy sequence patterns only antecedent matching of V_q is considered. This score emphasizes topic transitions learned from the global analysis over the data, which promotes videos that are consistent with frequent topic sequences over the entire corpus.

Local sequence score (LS). We retrieve M additional support and confidence score values from mined local sequential patterns with antecedent matching a subset of V_q and consequent matching a subset of V_r . The LS score is

$$S_{LS}(V_q, V_r) = \frac{1}{M} \sum_{i=0}^M c_i \frac{y_i}{|V_r|} + \frac{s_i}{D_q},$$

where D_q is total number of mined local sequences with antecedent values matching any subset of V_q , and y_i is the length of the matched subset of V_r . The LS score captures how specific topics from V_q relate to the other topics that appear in close proximity in the sequence database (local analysis).

We apply feature scaling to bring all scores into the range $[0, 1]$ before linearly fusing the scores with uniform weights in MOOCex, based on the results of pilot experiments. From the initial content-based recommendation, the above re-ranking technique offers more diverse videos in context using the knowledge of topic transitions (G2 and G3).

The final step in recommendation is aggregating the recommended videos to determine top- N prominent video sub-sequences. We scan videos from the top of the ranked list, and group them in sequential order if they are adjacent in a course syllabus. We constrain the length of sub-sequences to a maximum of four, and greedily search down on the list until there are N sub-sequences. Note that some sub-sequences may contain only one video. We re-order the sub-sequences based on the average ranking scores to form our final recommendation. Compared to traditional video recommendation, this method can provide more contextual information for learners to understand course concepts (G2).

Visual Interface

The MOOCex interface consists of three parts: a Video Panel, a Recommendation Panel, and a Configuration Panel (Figure 3). The *Video Panel* is a normal media player in which learners can watch a selected video. The *Recommendation Panel* is the main interface where a learner can explore recommended videos (or

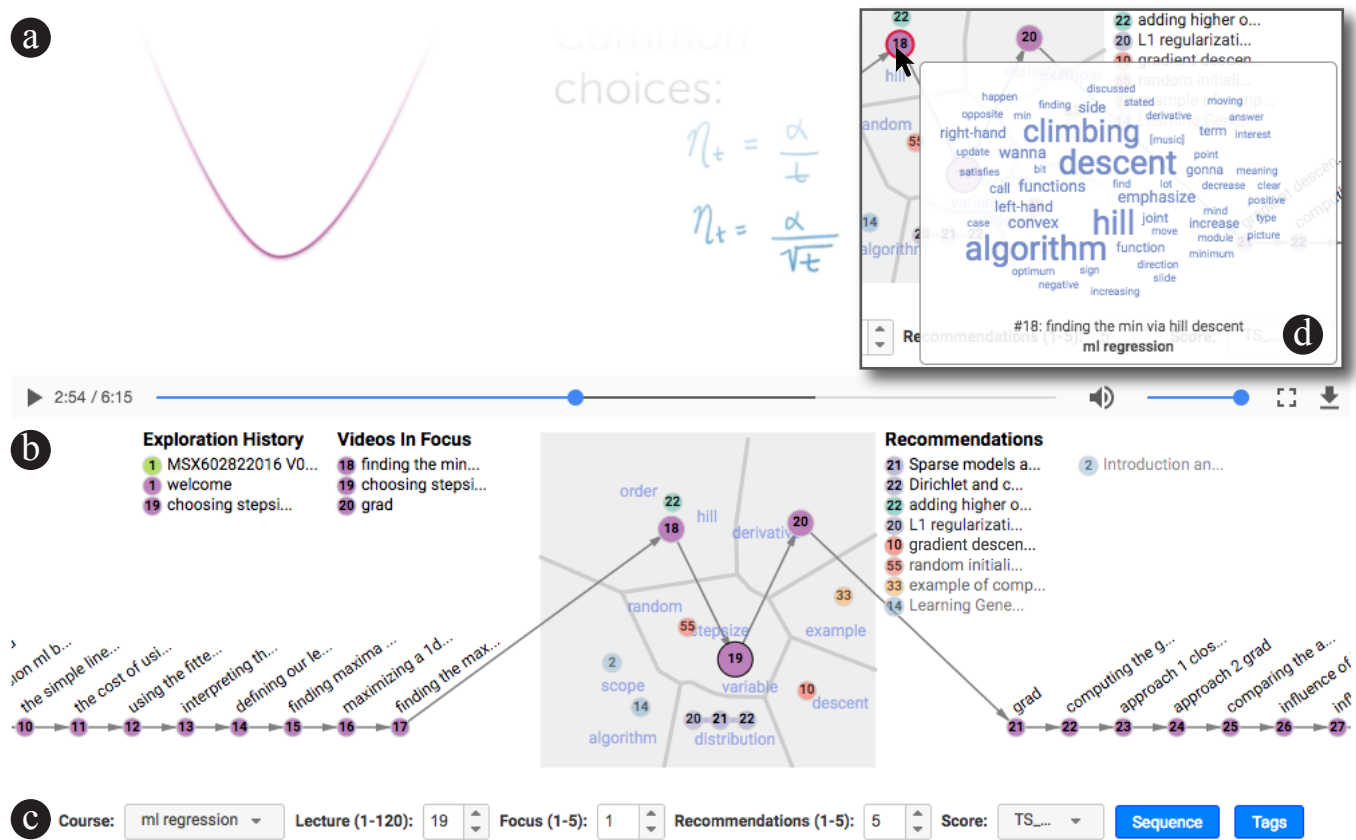


Figure 3. A learner is using MOOCex to explore recommendations of video (19) in *Machine Learning: Regression*. The interface consists of three main components: a) a Video Panel (cropped), b) a Recommendation Panel, and c) a Configuration Panel, d) Hovering over a video initiates a tooltip with a tag cloud to describe the content of the video.

video sub-sequences) and understand relationships between them, to inform their choice of an appropriate video to watch next. In the *Configuration Panel*, a learner can manipulate basic parameters about how recommendations are displayed, as well as select specific courses and lecture videos to view. Overall, we strive to design a visualization that is simple and easily used while conveying essential knowledge for learners to better determine their next step.

Visualization of Recommendations

The Recommendation Panel displays the current video, its neighboring videos, and the recommendations in a two-dimensional *Exploration Canvas* in the middle, and shows other videos in the current course linearly on both sides (Figure 3-b). Videos in the original course of the current video are connected with gray arrows to indicate their order in the syllabus. Each video is represented as a circle with a number indicating its position in that course syllabus. Color hues encode different courses. In the middle area for exploring recommendations, color opacity indicates the rank of that video in the recommendation list (the lighter the color, the lower the rank).

Unlike traditional ranked lists, we employ MDS [48] to position videos on the Exploration Canvas based on their distances in the recommendation space, the closer the more relevant. As in MDS projection only the relative distance

between items has meaning and the axes do not, we rotate the layout to make the general direction of videos in the current course flow from left to right, aligning with the natural sequence of other videos that are on either side. This rotation eases comprehension of the visualization because it matches people’s common sense. To obtain the angle to rotate, we compute the center of the videos before the current video in the MDS layout and that of those after, then form a vector from the previous center to the next, and use the angle between this vector and the positive direction of x-axis. Figures 4-a and -b show the rotated layouts of having two neighboring videos and four of them for (9) in *Machine Learning: Classification*. However, zig-zags in a longer video sequence could occur, which cannot be completely corrected. But learners usually focus more on semantics in a local space, thus not including too many neighboring videos in the Exploration Canvas.

Moreover, when performing MDS, we treat the recommended video sub-sequences as a unit (i.e., by averaging their recommendation distances to other videos), and then render videos of the sub-sequence from left to right in a roughly horizontal line and connect them together (e.g., 39-40 and 14-15-16 in Figure 4-a). To minimize overlap of circles, we later apply a repulsive force between videos to obtain the final layout. A learner can disable this sequence-based layout and position each video individually based on the MDS projection to gain a different perspective of the recommendations (Figure 4-c).




Figure 4. Layouts of recommendations and videos-in-focus in different situations when watching video 9 in *Machine Learning: Classification*; a) with two neighboring videos, b) with four neighboring videos, and c) with videos in sub-sequences separated.

To help learners better make sense of the MDS layout, we perform an agglomerative clustering [51] of the videos, and split the Exploration Canvas into multiple regions based on the clusters (Figure 4). Each region exhibits a relatively coherent set of topics and concepts in videos. We use the agglomerative approach for clustering because it does not require prior knowledge of the number of clusters. The boundaries of the regions are shown as subtle white polylines, determined by aggregating Voronoi cells of videos in the same cluster.

In addition, we overlay frequent topical keywords extracted from the text transcripts of each video cluster to reveal contextual information of different regions in the MDS projection (Figure 3). To obtain discriminative keywords, we first employ the standard TF-IDF method and then re-weight the keywords based on terms in video titles, because the titles are created by human thus providing high quality summarization. Next, we perform post-processing of the keywords based on the video clusters to remove duplications. These keywords are placed using a force-directed layout, and they can be hidden if users feel overwhelmed. For better positioning of the keywords, more advanced methods such as energy based optimization [20] can be investigated in the future. Our goal here is to provide some semantic structure (not necessarily precisely) in the information space, allowing everyday users to better understand the MDS layout.

In summary, the above visual design allows learners to better understand the semantics in the recommendation space based on the MDS layout, providing rich context of relevant concepts and videos, i.e., how they relate to each other in the current local space of learning (G2). Also, the recommendation and visualization of video sub-sequences offers additional contextual information with the recommended content (G2). These aspects form an informative and expressive view from which learners can better decide what to do next (G1).

Interactions and Miscellaneous Information

To facilitate exploration of videos, the Recommendation Panel displays auxiliary information on both sides (Figure 3-b). Videos that were recently visited and adjacent videos from the current course are shown in two vertical lists on the left. Similarly, recommended videos are shown on the right in a ranked list similar to the traditional approach. Interactive linking of the same video is provided as it is hovered over in the lists or in the Exploration Canvas. Meanwhile, a

tooltip pops-up showing a set of important keywords extracted from the video transcript and title based on the RAKE algorithm [42] (Figure 3-d). Also, clicking any of the videos selects it as the current video and updates the visualization. The above features allow learners to quickly get a basic sense of the videos and navigate through the information space.

Example Use Cases

Here, we demonstrate several possible usages of MOOCex when learners have different goals.

Suppose that Jack is learning *Machine Learning: Regression*, and he has just finished watching 19 titled with *Choosing stepsize and convergence criteria* (Figure 3). From his past experience, Jack already knows a bit about this topic, so now he wants to expand the scope of his knowledge on this aspect. On the Exploration Canvas, he finds a number of relevant videos that are close to the video he has just watched, including 55: *random initialization* and 10: *gradient descent intuition*, both in *Introduction to Machine Learning*, and a sub-sequence 20-21-22 in *Machine Learning (Undergraduate)*. As Jack wants to know more tricks for training a model, he decides to watch 55 next. Clicking that video sets it as the current video, and brings him even more related content. Jack could keep diving into it to get more knowledge or just simply go back to the original course.

Now imagine that Mary is also learning this course and has just finished watching the same video (Figure 3). However, she does not quite understand it. By looking at the recommendations, Mary finds 10: *gradient descent intuition* is close to the current video 19 with a title and tooltip keywords suggesting it has helpful background, but is positioned in a different region indicating a slightly different perspective. Such a recommendation with high semantic similarity depicted by its relative proximity can provide complementary information for better understanding of the current video.

Finally, consider that another user, Mike, thinks he fully grasps concepts in all the videos that he has watched in this course. He does not want to explore other topics related to the current video 19, because he wants to continue learning. After glancing the title of the next video 20: *Gradients derivatives in multiple dimension*, Mike happens to know this topic to

some extent from other resources. But he discovers that there is a video close to it, 39: *Example of computing derivative for logistic regression in Machine Learning: Classification*, which could broaden his view on this topic and better prepare him for the next video. So he decides to watch this video first before moving forward.

In summary, the diversity of recommendations offers learners opportunities and flexibility to gain specific knowledge based on their different learning objectives and past experiences (G3). However, if the recommendations are displayed as a linear ranked list, a learner probably just visits a few items near the top of the list. The above scenarios demonstrate that this visualization which provides the context and semantics of recommendations could lead to a deeper and more thorough understanding of the information space (G2) and thus guide learners to make a better decision (G1).

EVALUATION

We carried out three experiments to evaluate the effectiveness and usefulness of MOOCex. The first was conducted to quantitatively assess the quality of results provided by the recommendation engine. The purposes of the second and the third experiments were to understand how MOOCex can be used on MOOC platforms in a qualitative manner, from both MOOC instructor and learner perspectives.

Evaluation of Recommendation

Guided by our design goal G3, we aimed to measure the diversity of recommended results and compare the sequence-based recommendation with the traditional content-based approach (CB). The content-based method utilizes text transcripts of videos without considering the sequential relationships of topics and concepts in courses. We observed that the results of CB concentrate within the same course of V_q , which may be possibly based on the distinctiveness of each instructor’s language usage. To provide more diverse recommendations, we filtered out videos from the same course for this evaluation, thus better assessing the ability of CB and our method for providing users with flexibility in learning.

Metrics

We employed the following two metrics to quantitatively evaluate the quality of recommendations:

TR: Topic redundancy measures the average number of repeated items in top- N recommendations between corresponding query videos from the same section of a syllabus. The section groupings in course syllabi created by instructors indicate closely related concepts. The intuition here is that as learners consume these videos, recommendations should enable further opportunities for enriched understanding (i.e., lower redundancy).

CD: Course diversity measures the average number of distinct courses from which the top- N recommendations originate. Intuitively, we want the recommended videos to span across different courses (i.e., higher diversity), which may provide a more flexible learning experience.

Perhaps a more general metric for the recommender performance is accuracy (or precision and recall), which is used

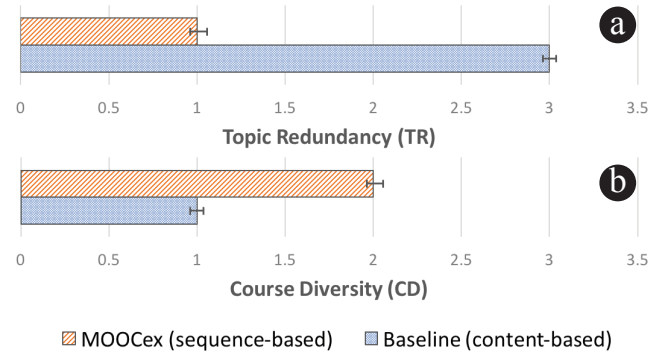


Figure 5. Comparison of recommended results from MOOCex (sequence-based) and baseline (content-based) on: a) the averages of topic redundancy (shorter bars are better), and b) course diversity (longer bars are better). Error bars show 95% CI.

in many previous work [6, 22]. However, this requires the ground truth: learners’ actual behaviors in selecting videos from different courses. But this real-world information is not available, because most MOOC platforms do not adequately support this interaction and a cross-platform video navigation system does not exist yet. Therefore, in this paper, we evaluate the diversity of recommendation guided by G3.

Results

We computed the above two metrics with $N = 10$ for all videos in our dataset (4,186 videos). Figure 5 shows the results that compare the recommendation engine of MOOCex (sequence-based) and the baseline (content-based). We can observe that MOOCex performs better in both metrics (TR: $\mu = 1$, CI [0.91, 1.09]; CD: $\mu = 2$, CI [1.94, 2.06]), compared to those in baseline (TR: $\mu = 3$, CI [2.88, 3.12]; CD: $\mu = 1$, [0.96, 1.04]). Further, the non-overlapping 95% CIs indicate that these effects are significant. These results indicate that MOOCex can provide recommendations with more diverse perspectives.

Interviews with MOOC Instructors

We conducted semi-structured interviews with MOOC instructors to collect qualitative feedback about MOOCex.

Experimental Setup

We recruited two MOOC instructors from different universities. They both have many years of teaching experience in traditional classrooms and have taught several MOOCs in recently. One was from a computer science background (E1), and the other was specialized in quantitative methods in social science (E2). During the interviews, we first introduced some background and the features of MOOCex, and then asked them to try the tool, during which we collected their comments. Think aloud protocol was used, and we required them to give feedback from both instructors’ and students’ perspectives. We recorded the entire sessions and took notes when necessary. Each interview lasted about a hour.

Results

In general, the instructors appreciated the tool and valued its potential for enhanced flexibility for learning. They both agreed that the current fixed course syllabi have big limitations for certain learner groups. Also, they were eager to apply the tool in their MOOC offerings and were curious about how it could affect students’ performance and learning behaviors.

Students' perspectives. Thinking from a student's point of view, the instructors both agreed that the tool could broaden their vision and deepen their understanding of course materials, especially for professional learners. They also thought that the interactivity of the tool could help students become more engaged with the course. E1 mentioned that they usually provided recommendations of other materials in the end of a course to give students future learning directions. “*Students can have more freedom and gain richer context with this kind of per-video recommendation,*” she further added. E2 particularly liked the visual representation of videos on the Exploration Canvas. He said that the keywords, the regions, and the positions of videos offered “*a holistic image of what students are learning currently,*” which could be helpful for them to grasp the concepts of a course and its progress. He also commented that the connected subsequences of videos in the visualization could provide students with a good awareness of the context of knowledge or related concepts.

Instructors' perspectives. The instructors were excited about the capabilities of the tool for improving their teaching. They said that MOOCex could be very useful for course preparation. E1 commented: “*I normally don't look at what others teach, but the tool provides the awareness of related lectures, so I could borrow some materials to enhance my lecture, and avoid unnecessary duplication.*” Both E1 and E2 emphasized that MOOCex could be used for dynamically guiding the support offered on course forums, for example, pointing out details and questions covered in recommended videos but not in current course. E2 commented that the visualization could provide objective feedback to the design of a course. For example, he said: “*If you see one lecture is here [on the Exploration Canvas], then you go very far for the second lecture, and back here again for the third lecture, you should really think about reordering the content presented in the videos.*”

Problems and issues. From the interview, E1 commented that the visualization might be confused with a concept (knowledge) map. She explained that it could be because arrows only exist between videos of the current course, and suggested also adding dependency between recommendations. For the current MOOCex design, we only included sequential information of consecutive videos. Both E1 and E2 were worried that students might think: in order to fully grasp the knowledge, watching videos within a whole region on the visualization is needed. They were also afraid that students could dive too deep through recommendations but not move forward. Moreover, they both noticed that sometimes the text labels were not properly positioned, such as overlapping with each other and going across region boundaries.

General comments. Overall, the instructors thought highly about MOOCex and wanted to use it in their MOOC offerings in the future. They also provided several suggestions for enhancing the tool. One was to accommodate the recommendation of sub-segments of videos, because some MOOC lectures may cover multiple key concepts in relatively long videos. Moreover, E2 suggested dynamically loading videos of similar topics by clicking keywords on the Exploration Canvas. He further demanded a more adaptive recommendation engine




Figure 6. Recommendation Panel in the baseline condition, which only presents the recommendation in a list.

that suggests videos differently based on the number of videos in the current course on the canvas. Also, E1 pointed out that it would be interesting to apply MOOCex in non-MOOC video recommendation, for example, in corporate training.

Laboratory Study with MOOC Learners

We further carried out a laboratory study to better understand how MOOCex can be used by MOOC learners.

Experimental Setup

We recruited 12 participants (9 males, 3 females; aged 20–50) from an IT company; all of them have taken MOOCs before and hold Masters and/or PhDs in computer science or a related field, which represents the professional learner group for MOOCex. They all had some level of knowledge about machine learning (but not experts) to match our video data in the experiment being machine learning related. This allowed participants to better make sense of the content. Further, all of them had experienced video recommendations on other sites like YouTube and Netflix and all had taken MOOCs in the past (but might not complete them).

We compared MOOCex with a baseline view of the recommendations in a list by hiding the Exploration Canvas in Figure 6. The experimental task was to select the most meaningful video to watch next from the recommendations of a particular video, with the presentation in one of the two conditions. Participants were asked to talk aloud why they chose such video, and then watch/skim that video to indicate if they felt it was a good choice. We understand that the choice may be affected by many subjective matters. However, we are more interested in behavioral changes of participants across the two conditions, because there is no right answer for choosing videos. We hypothesize that participants have more deliberate reasoning in mind for taking actions in the MOOCex condition.

A within-subject design was employed. We sampled six videos (2–4 minutes length each) from our dataset, all related to machine learning concepts, to form six tasks for the study. The tasks were randomly associated with the two conditions (three each) for every participant. During the study, the two conditions (MOOCex and baseline) were presented to participants in a counter-balanced order. Within each condition, they first received a demonstration of the tool, during which they could ask questions. Then, participants completed one trial task and three actual tasks. The trial task videos were different from the actual ones, but stayed the same across participants. Participants were instructed to either skim or watch videos during the tasks. Finally, participants filled

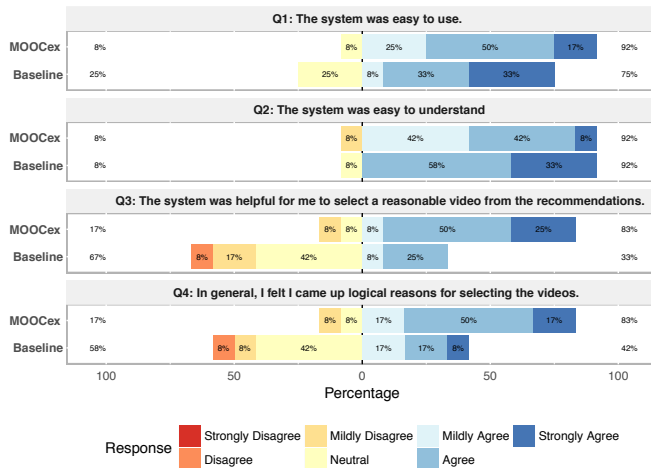


Figure 7. Comparison of participants' agreements on MOOCex and baseline from the post-study questionnaire (with a 7-point Likert scale). Agree scores to the right, Neutral or Disagree to the left. No participant selected Strongly Disagree.

in a post-study questionnaire and received a short informal interview. The study lasted approximately 40 minutes.

Results

During the experiment, we recorded the time participants spent selecting a video in each task. On average, participants took a bit more time with MOOCex: $\mu = 123.8$ seconds, CI [99.0, 148.7], than baseline: $\mu = 118.6$ seconds, CI [97.0, 140.2]. This is reasonable because participants had to cognitively process more information with MOOCex. However, the large overlapping confidence intervals indicate that the effect was not significant. From the responses in the post-study questionnaire, participants felt they had made a good and logical choice in 32 out of 36 tasks (88.9%) with MOOCex, compared to 25 out of 36 tasks (69.4%) with baseline. Although subjective, these results indicate that MOOCex helped participants make better and more confident choices.

The results of the post-study questionnaire further show that participants thought MOOCex was roughly as easy to use and about as easy to understand as baseline (Figure 7). Although there were more visual elements in MOOCex than baseline, participants did not find it more difficult to use (MOOCex: $M = 6$, $IQR = 1$; baseline: $M = 6$, $IQR = 2.25$). MOOCex was perceived slightly less easier to understand, which is plausible because it was a new interface. However, the effect was small (MOOCex: $M = 5.5$, $IQR = 1$; baseline: $M = 6$, $IQR = 1$), indicating that participants accepted MOOCex well and quickly.

As shown in Figure 7, we also observe that participants thought MOOCex was more helpful in selecting videos from recommendations (MOOCex: $M = 6$, $IQR = 0.5$; baseline: $M = 4$, $IQR = 1.5$), and helped them make more logical choices (MOOCex: $M = 6$, $IQR = 1$; baseline: $M = 4$, $IQR = 1.25$). This indicates MOOCex is more effective for users to make sense of the information space and select reasonable videos from recommendations.

In terms of the list ranking, of all 36 tasks, 8 selections (22.2%) were made to the top (first) recommended video with MOOCex

and 13 (36.1%) for those with baseline. Picking within the top-5 recommendations, 31 selections (86.1%) comprised the MOOCex condition compared to 25 (69.4%) in the baseline. In other words, when presented with a list, participants typically clicked the first item (which is to be expected). When using MOOCex, participants typically clicked one of the top-5 results. From records in the study, we find in the baseline that 32.0% (8 out of 25) of those top-5 choices did not match what they wanted (i.e., regret rate). For MOOCex, the top-5 choice regret rate was much lower (6.4%; 2 out of 31). When picking outside the top-5 in the baseline, the regret rate was 27.3% (3 out of 11) versus 40.0% (2 out of 5) in MOOCex.

Moreover, we found that participants had strategy changes for selecting videos from recommendations under the two conditions. With baseline, almost all participants scanned from the top of the ranked lists to select videos. Some thought more about the topics of the current video or the next one to make their decisions, while others tended to pick videos that "look interesting." Participants also indicated that the tooltips were more helpful in the baseline condition. With MOOCex, participants usually searched from the near to far neighbors of the current video on the Exploration Canvas. They also stated that the background regions and keywords had affected their behaviors. For example, three participants mentioned using a top-down approach for video selection, and many said they made comparisons of candidate videos in different regions. Two participants mentioned that the connected sub-sequences of videos helped them find where to start if they were interested about the topics. There was one participant who generally used the same list-scanning strategy in both conditions citing "The map [Exploration Canvas] is too busy." He did however indicate the visualization of MOOCex did influence some of his choices because he also observed the positions of the videos when he hovered over it in the ranked list.

DISCUSSION

We have presented a visual recommender system, MOOCex, designed to recommend lectures from a set of courses to provide multiple perspectives and supplements to an enrolled course. MOOCex features a spatial-visual, semantic representation to allow learners to understand a recommended video's topic in relation to their current context in an enrolled course. The visualization requires a recommendation engine based on sequential inter-topic relationships, which provides more exploration with the top-5 highly ranked recommendations (instead of just promoting the top ranked video). Here, we discuss the limitations of MOOCex, implications obtained from our studies, and the generalization of the technique.

Limitations

Although effective, MOOCex is not without limitations. First, we utilize sequential relationships between videos to provide higher-quality and more diverse recommendations, but the hierarchies of MOOC syllabi are not considered in this paper. This information could be helpful for offering recommendations at multiple levels of concepts. Second, in our current layout, some keywords may overlap with other elements and longer keywords or sub-sequences may cross the boundary of a Voronoi cell. Alternative layout methods [20]

and visual aggregation of video sub-sequences can be applied to resolve the issues. Third, our current summarization (on the tooltip) of a video's content is purely based the text transcripts. Some visual cues such as keyframes extracted from the videos could be more informative. Similarly, videos now encoded as circles could be a bit abstract to users. Using thumbnails might allow for more effective browsing of the videos, however, this may result in visual clutter as each video requires more space to display in the visualization.

Ideally, a longitudinal deployment study is needed to provide deeper insights on how MOOCex is used in practice, and find out whether the learning outcomes are improved or not, which is our ultimate goal. However, current popular MOOC platforms do not easily allow for such integration. We argue that the three experiments conducted to assess MOOCex from different perspectives already provide us with a relatively thorough picture of the tool's performance. A deployment study is left for future work.

Our evaluation of the recommendation engine is also limited and heuristic. Due to the lack of ground truth of user behaviors in navigating videos of different courses in a cross-platform system, metrics such as accuracy, precision, and recall cannot be applied, which is an obvious weakness of the evaluation. The main obstacle is that such data is not self-contained within an actual existing system, thus difficult to capture. We believe a future deployment of MOOCex to a real audience will enable us to collect a large number of relevant user sessions for the evaluation. This also opens new opportunities to further improve our recommender algorithm based on user browsing history, such as applying collaborative filtering [52].

Design Implications

Our study indicates several important considerations for the future design of interactive visual recommender systems for MOOCs. Ultimately, the goals of providing recommendations are different between MOOC platforms and common video browsing websites such as YouTube. While the latter has playlists and genres, MOOC content conveys educational knowledge information contained in the course content structure. The goal is to help users obtain this diverse knowledge and learn more effectively (as opposed to having viewers stay on site watching content and advertisements). For the MOOC learning scenario, there exist two conflicting design goals: encouraging learners to watch more videos to broaden their perspectives, and preventing them from diving too deep in the unlimited space of videos. Finding a balance between the two is a critical consideration for system design.

Additionally, our study reveals the importance of sequential inter-topic relationships in providing better recommendations of MOOC videos as well as the importance of explicitly visualizing such sequentiality. Hence a proper ranking of recommendations must be developed to obtain such a sequence across the corpus. This aligns with some of the observations from earlier work [25, 44], where sequential dependency is extended to a graph. Showing this information is critical, but is only one aspect. Another dimension is the semantic relationship between videos, which is what we also strive to present in MOOCex (e.g., keywords, and regions on the

Exploration Canvas). However, it is not easy to generate a visualization that conveys well these two types of information simultaneously, which is another essential aspect to consider in the design.

Generalization

The MOOCex technique, or part of it, is generalizable, although we design and develop it in the context of MOOC learning. For example, we can extend the recommendation engine by computing the signatures of videos with additional information, such as audio and image frames, instead of merely the topics of transcripts as in the current implementation. Also, the front-end interface of MOOCex is not constrained with the sequence-based recommendation back-end, which can be applied for exploring the recommendations of more generic video browsing platforms.

Further, although we currently focus on whole video analysis, the sequence-based approach and the interactive visualization can be generalized to video sub-segment recommendations. This could be especially helpful when a MOOC video is longer or covers multiple learning concepts. A preprocessing step would be splitting original videos into sub-segments, each consisting of one specific low-level concept.

CONCLUSION AND FUTURE WORK

We have presented MOOCex, an interactive tool that offers a flexible learning experience with dynamic recommendations and visual exploration of MOOC videos. Its recommendation engine considers sequential relationships between videos based on course syllabi, to ease the learning of concepts. Also, its visual interface supports a richer semantic representation of videos in the information space, allowing learners to quickly make sense of the recommendations and decide their next step. Three studies are conducted to evaluate different aspects of MOOCex by comparing it with traditional methods (i.e., content-based recommendation and ranked list), and the results indicate that MOOCex is effective and useful in various MOOC learning scenarios.

In the future, we plan to extend our recommendation engine to including more video features such as audio and image frames, and further enhance the visual interface. We also would like to experiment with recommending and exploring video sub-segments, as well as non-MOOC scenarios. Finally, we aim to conduct a deployment study to collect real-world user data, thus further evaluating both the recommender algorithm and the visual interface of MOOCex.

ACKNOWLEDGMENT

We wish to thank Dr. Haiyi Zhu and Dr. Hao Dong for their valuable comments on the system, and the users who participated in our studies. We also want to thank all the FXPAL staff for their support of this project.

REFERENCES

1. J. Adcock, M. Cooper, L. Denoue, H. Pirsiavash, and L. A. Rowe. 2010. TalkMiner: A Lecture Webcast Search Engine. In *Proceedings of the ACM International Conference on Multimedia (MM '10)*. 241–250. DOI: <http://dx.doi.org/10.1145/1873951.1873986>

2. R. Agrawal, M. Christoforaki, S. Gollapudi, A. Kannan, K. Kenthapadi, and A. Swaminathan. 2014. *Mining Videos from the Web for Electronic Textbooks*. Springer International Publishing, 219–234. DOI: http://dx.doi.org/10.1007/978-3-319-07248-7_16
3. B. Bach, C. Shi, N. Heulot, T. Madhyastha, T. Grabowski, and P. Dragicevic. 2016. Time Curves: Folding Time to Visualize Patterns of Temporal Evolution in Data. *IEEE Transactions on Visualization and Computer Graphics* 22, 1, 559–568. DOI: <http://dx.doi.org/10.1109/TVCG.2015.2467851>
4. D. M. Blei, A. Y. Ng, and M. I. Jordan. 2003. Latent dirichlet allocation. *Journal of Machine Learning Research* 3, 993–1022.
5. D. H. Chau, A. Kittur, J. I. Hong, and C. Faloutsos. 2011. Apolo: Making Sense of Large Network Data by Combining Rich User Interaction and Machine Learning. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. 167–176. DOI: <http://dx.doi.org/10.1145/1978942.1978967>
6. K. Choi, D. Yoo, G. Kim, and Y. Suh. 2012. A Hybrid Online-product Recommendation System: Combining Implicit Rating-based Collaborative Filtering and Sequential Pattern Analysis. *Electronic Commerce Research and Applications* 11, 4, 309–317. DOI: <http://dx.doi.org/10.1016/j.elerap.2012.02.004>
7. J. Choo, C. Lee, C. K. Reddy, and H. Park. 2013. UTOPIAN: User-Driven Topic Modeling Based on Interactive Nonnegative Matrix Factorization. *IEEE Transactions on Visualization and Computer Graphics* 19, 12, 1992–2001. DOI: <http://dx.doi.org/10.1109/TVCG.2013.212>
8. E. Clarkson, K. Desai, and J. Foley. 2009. ResultMaps: Visualization for Search Interfaces. *IEEE Transactions on Visualization and Computer Graphics* 15, 6, 1057–1064. DOI: <http://dx.doi.org/10.1109/TVCG.2009.176>
9. C. Coffrin, L. Corrin, P. de Barba, and G. Kennedy. 2014. Visualizing Patterns of Student Engagement and Performance in MOOCs. In *Proceedings of the Fourth International Conference on Learning Analytics And Knowledge (LAK '14)*. 83–92. DOI: <http://dx.doi.org/10.1145/2567574.2567586>
10. F. Dernoncourt, C. Taylor, U.-M. O'Reilly, K. Veeramachaneni, S. Wu, C. Do, and S. Halawa. 2013. MoocViz: A Large Scale, Open Access, Collaborative, Data Analytics Platform for MOOCs. In *Proceedings of NIPS Workshop on Data-Driven Education*.
11. C. di Sciascio, V. Sabol, and E. E. Veas. 2016. Rank As You Go: User-Driven Exploration of Search Results. In *Proceedings of the 21st International Conference on Intelligent User Interfaces (IUI '16)*. 118–129. DOI: <http://dx.doi.org/10.1145/2856767.2856797>
12. M. Dörk, S. Carpendale, and C. Williamson. 2011. The Information Flaneur: A Fresh Look at Information Seeking. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. 1215–1224. DOI: <http://dx.doi.org/10.1145/1978942.1979124>
13. M. Dörk, R. Comber, and M. Dade-Robertson. 2014. Monadic Exploration: Seeing the Whole Through Its Parts. In *Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems (CHI '14)*. 1535–1544. DOI: <http://dx.doi.org/10.1145/2556288.2557083>
14. M. Dörk, N. H. Riche, G. Ramos, and S. Dumais. 2012. PivotPaths: Strolling through Faceted Information Spaces. *IEEE Transactions on Visualization and Computer Graphics* 18, 12, 2709–2718. DOI: <http://dx.doi.org/10.1109/TVCG.2012.252>
15. S. Doroudi, K. Holstein, V. Alevan, and E. Brunskill. 2016. Sequence Matters, But How Exactly? A Method for Evaluating Activity Sequences from Data. In *Proceedings of the International Conference on Educational Data Mining*. 70–77.
16. D. M. Edwards and L. Hardman. 1999. Lost in Hyperspace: Cognitive Mapping and Navigation in a Hypertext Environment. In *Hypertext: Theory into Practice*, Ray McAleese (Ed.). Intellect Books, 90–105. <http://dl.acm.org/citation.cfm?id=320628.320639>
17. P. Fournier-Viger, A. Gomariz, T. Gueniche, E. Mwamikazi, and R. Thomas. 2013. TKS: Efficient Mining of Top-K Sequential Patterns. In *Advanced Data Mining and Applications*, Hiroshi Motoda, Zhaohui Wu, Longbing Cao, Osmar Zaiane, Min Yao, and Wei Wang (Eds.). Springer Berlin Heidelberg, 109–120. DOI: http://dx.doi.org/10.1007/978-3-642-53914-5_10
18. P. Fournier-Viger and V. S. Tseng. 2013. TNS: Mining Top-k Non-redundant Sequential Rules. In *Proceedings of the 28th Annual ACM Symposium on Applied Computing (SAC '13)*. 164–166. DOI: <http://dx.doi.org/10.1145/2480362.2480395>
19. S. Fu, J. Zhao, W. Cui, and H. Qu. 2017. Visual Analysis of MOOC Forums with iForum. *IEEE Transactions on Visualization and Computer Graphics* 23, 1, 201–210. DOI: <http://dx.doi.org/10.1109/TVCG.2016.2598444>
20. E. Gomez-Nieto, F. S. Roman, P. Pagliosa, W. Casaca, E. S. Helou, M. C. F. de Oliveira, and L. G. Nonato. 2014. Similarity Preserving Snippet-Based Visualization of Web Search Results. *IEEE Transactions on Visualization and Computer Graphics* 20, 3, 457–470. DOI: <http://dx.doi.org/10.1109/TVCG.2013.242>
21. S. Gratzl, A. Lex, N. Gehlenborg, H. Pfister, and M. Streit. 2013. LineUp: Visual Analysis of Multi-Attribute Rankings. *IEEE Transactions on Visualization and Computer Graphics* 19, 12, 2277–2286. DOI: <http://dx.doi.org/10.1109/TVCG.2013.173>

22. N. Hariri, B. Mobasher, and R. Burke. 2012. Context-aware Music Recommendation Based on Latenttopic Sequential Patterns. In *Proceedings of ACM Conference on Recommender Systems (RecSys '12)*. 131–138. DOI: <http://dx.doi.org/10.1145/2365952.2365979>
23. M. A. Hearst. 1995. TileBars: Visualization of Term Distribution Information in Full Text Information Access. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '95)*. 59–66. DOI: <http://dx.doi.org/10.1145/223904.223912>
24. S. Jiang, X. Qian, T. Mei, and Y. Fu. 2016. Personalized Travel Sequence Recommendation on Multi-Source Big Social Media. *IEEE Transactions on Big Data* 2, 1, 43–56. DOI: <http://dx.doi.org/10.1109/TBDATA.2016.2541160>
25. KhanAcademy Knowledge Map. Accessed in 2017. <https://www.khanacademy.org/exercisedashboard>.
26. H. Kim, J. Choo, H. Park, and A. Endert. 2016. InterAxis: Steering Scatterplot Axes via Observation-Level Interaction. *IEEE Transactions on Visualization and Computer Graphics* 22, 1, 131–140. DOI: <http://dx.doi.org/10.1109/TVCG.2015.2467615>
27. J. Kim, P. J. Guo, D. T. Seaton, P. Mitros, K. Z. Gajos, and R. C. Miller. 2014. Understanding In-video Dropouts and Interaction Peaks Inonline Lecture Videos. In *Proceedings of the First ACM Conference on Learning @ Scale Conference (L@S '14)*. 31–40. DOI: <http://dx.doi.org/10.1145/2556325.2566237>
28. M. Kim, K. Kang, D. Park, J. Choo, and N. Elmqvist. 2017. TopicLens: Efficient Multi-Level Visual Topic Exploration of Large-Scale Document Collections. *IEEE Transactions on Visualization and Computer Graphics* 23, 1, 151–160. DOI: <http://dx.doi.org/10.1109/TVCG.2016.2598445>
29. J. S. Kinnebrew, K. M. Loretz, and G. Biswas. 2013. A Contextualized, Differential Sequence Mining Method to Derive Students' Learning Behavior Patterns. *Journal of Educational Data Mining* 5, 1, 190–219.
30. Y. Kleiman, J. Lanir, D. Danon, Y. Felberbaum, and D. Cohen-Or. 2015. DynamicMaps: Similarity-based Browsing Through a Massive Set of Images. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. 995–1004. DOI: <http://dx.doi.org/10.1145/2702123.2702224>
31. K. Klouche, T. Ruotsalo, D. Cabral, S. Andolina, A. Bellucci, and G. Jacucci. 2015. Designing for Exploratory Search on Touch Devices. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. 4189–4198. DOI: <http://dx.doi.org/10.1145/2702123.2702489>
32. B. Y.-L. Kuo, T. Hentrich, B. M. . Good, and M. D. Wilkinson. 2007. Tag Clouds for Summarizing Web Search Results. In *Proceedings of the 16th International Conference on World Wide Web (WWW '07)*. 1203–1204. DOI: <http://dx.doi.org/10.1145/1242572.1242766>
33. C. D. Manning, P. Raghavan, and H. Schütze. 2008. *Introduction to Information Retrieval*. Cambridge University Press.
34. T. Mei, B. Yang, X.-S. Hua, and S. Li. 2011. Contextual Video Recommendation by Multimodal Relevance and User Feedback. *ACM Transactions on Information Systems* 29, 2, 24 pages. DOI: <http://dx.doi.org/10.1145/1961209.1961213>
35. C. R. Morales, A. P. Pérez, S. V. Soto, C. H. Martinez, and A. Zafra. 2006. Using sequential pattern mining for links recommendation in adaptive hypermedia educational systems. In *Current Developments in Technology Assisted Education*, Vol. 2. 1016–1020.
36. T. Nguyen and J. Zhang. 2006. A Novel Visualization Model for Web Search Results. *IEEE Transactions on Visualization and Computer Graphics* 12, 5, 981–988. DOI: <http://dx.doi.org/10.1109/TVCG.2006.111>
37. J. D. Novak and D. B. Gowin. 1984. *Learning how to learn*. Cambridge University Press.
38. D. Onah, J. Sinclair, and R. Boyatt. 2014. Dropout Rates Of Massive Open Online Courses: Behavioural Patterns. In *Proceedings of International Conference on Education and New Learning Technologies*. 5825–5834.
39. D. Parra, P. Brusilovsky, and C. Trattner. 2014. See What You Want to See: Visual User-driven Approach for Hybrid Recommendation. In *Proceedings of the 19th International Conference on Intelligent User Interfaces (IUI '14)*. 235–240. DOI: <http://dx.doi.org/10.1145/2557500.2557542>
40. J. Peltonen, K. Belorustceva, and T. Ruotsalo. 2017. Topic-Relevance Map: Visualization for Improving Search Result Comprehension. In *Proceedings of the 22Nd International Conference on Intelligent User Interfaces (IUI '17)*. 611–622. DOI: <http://dx.doi.org/10.1145/3025171.3025223>
41. H. Qu and Q. Chen. 2015. Visual Analytics for MOOC Data. *IEEE Computer Graphics and Applications* 35, 6, 69–75. DOI: <http://dx.doi.org/10.1109/MCG.2015.137>
42. S. Rose, D. Engel, N. Cramer, and W. Cowley. 2010. Automatic keyword extraction from individual documents. *Text Mining: Applications and Theory*, 1–20.
43. G. Salton, A. Wong, and C.-S. Yang. 1975. A vector space model for automatic indexing. *Commun. ACM* 18, 11, 613–620.
44. M. Schwab, H. Strobel, J. Tompkin, C. Fredericks, C. Huff, D. Higgins, A. Strezhev, M. Komisarchik, G. King, and H. Pfister. 2017. booc.io: An Education System with Hierarchical Concept Maps and Dynamic Non-linear Learning Plans. *IEEE Transactions on Visualization and Computer Graphics* 23, 1, 571–580. DOI: <http://dx.doi.org/10.1109/TVCG.2016.2598518>

45. C. Shi, S. Fu, Q. Chen, and H. Qu. 2015. VisMOOC: Visualizing video clickstream data from Massive Open Online Courses. In *Proceedings of IEEE Pacific Visualization Symposium (PacificVis)*. 159–166. DOI: <http://dx.doi.org/10.1109/PACIFICVIS.2015.7156373>
46. R. Spiro and J. Jehng. 1990. Cognitive flexibility and hypertext: Theory and technology for the nonlinear and multidimensional traversal of complex subject matter. In *Cognition, education and multimedia: Exploring ideas in high technology*, D. Nix and R. Spiro (Eds.). Lawrence Erlbaum.
47. A. Spoerri. 2004. RankSpiral: Toward Enhancing Search Results Visualizations. In *IEEE Symposium on Information Visualization*. 18–18. DOI: <http://dx.doi.org/10.1109/INFVIS.2004.56>
48. W. S. Torgerson. 1952. Multidimensional scaling: I. Theory and method. *Psychometrika* 17, 4, 401–419. DOI: <http://dx.doi.org/10.1007/BF02288916>
49. K. Verbert, D. Parra, P. Brusilovsky, and E. Duval. 2013. Visualizing Recommendations to Support Exploration, Transparency and Controllability. In *Proceedings of the International Conference on Intelligent User Interfaces (IUI '13)*. 351–362. DOI: <http://dx.doi.org/10.1145/2449396.2449442>
50. J. A. Wise. 1999. The ecological approach to text visualization. *Journal of the American Society for Information Science* 50, 13, 1224–1233. DOI: [http://dx.doi.org/10.1002/\(SICI\)1097-4571\(1999\)50:13<1224::AID-ASIS>3.0.CO;2-4](http://dx.doi.org/10.1002/(SICI)1097-4571(1999)50:13<1224::AID-ASIS>3.0.CO;2-4)
51. I. H. Witten, E. Frank, and M. A. Hall. 2011. *Data Mining: Practical Machine Learning Tools and Techniques* (3rd ed.). Morgan Kaufmann.
52. F. Zhang, K. Zheng, N. J. Yuan, X. Xie, E. Chen, and X. Zhou. 2015. A Novelty-Seeking Based Dining Recommender System. In *Proceedings of the 24th International Conference on World Wide Web (WWW '15)*. 1362–1372. DOI: <http://dx.doi.org/10.1145/2736277.2741095>
53. J. Zhao, C. Collins, F. Chevalier, and R. Balakrishnan. 2013. Interactive Exploration of Implicit and Explicit Relations in Faceted Datasets. *IEEE Transactions on Visualization and Computer Graphics* 19, 12, 2080–2089. DOI: <http://dx.doi.org/10.1109/TVCG.2013.167>
54. S. Zheng, M. B. Rosson, P. C. Shih, and J. M. Carroll. 2015. Understanding Student Motivation, Behaviors and Perceptions in MOOCs. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '15)*. 1882–1895. DOI: <http://dx.doi.org/10.1145/2675133.2675217>
55. Q. Zhu, M. L. Shyu, and H. Wang. 2013. VideoTopic: Content-Based Video Recommendation Using a Topic Model. In *Proceedings of the IEEE International Symposium on Multimedia*. 219–222. DOI: <http://dx.doi.org/10.1109/ISM.2013.41>