

# booc.io: An Education System with Hierarchical Concept Maps and Dynamic Non-linear Learning Plans

Michail Schwab, Hendrik Strobel, James Tompkin, Colin Fredericks, Connor Huff, Dana Higgins, Anton Strezhnev, Mayya Komisarchik, Gary King, Hanspeter Pfister

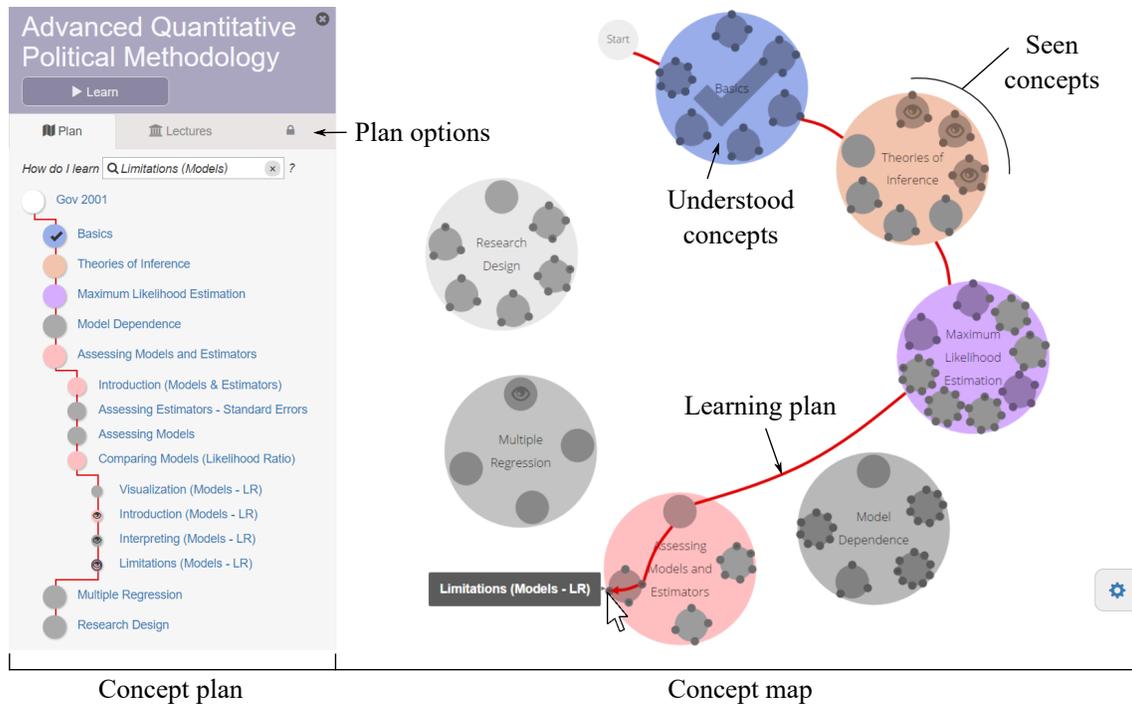


Fig. 1. The booc.io organizational structure. Concepts and their corresponding material are represented as circles arranged into a zoomable hierarchy—the *concept map*. Linear paths through the material are ordered around the clock face, with non-linear paths short-cutting inside. The structure is included as a linear list—the *concept plan*—in a left-hand-side panel. *In red*: to guide students to concepts, a personal *learning plan* automatically considers concept dependencies and material which is already seen or understood.

**Abstract**— Information hierarchies are difficult to express when real-world space or time constraints force traversing the hierarchy in linear presentations, such as in educational books and classroom courses. We present *booc.io*, which allows linear *and* non-linear presentation and navigation of educational concepts and material. To support a breadth of material for each concept, booc.io is Web based, which allows adding material such as lecture slides, book chapters, videos, and LTIs. A visual interface assists the creation of the needed hierarchical structures. The goals of our system were formed in expert interviews, and we explain how our design meets these goals. We adapt a real-world course into booc.io, and perform introductory qualitative evaluation with students.

**Index Terms**—Hierarchies, information visualization, education.

## 1 INTRODUCTION

Education is often supported by online Web components: to present rich material, to teach at a distance, to complement an in-person course with an online repository of supporting materials, or to provide sup-

plemental resources via search tools. To cope with these different use cases, modern online education platforms have many complex goals. They organize materials, provide access to different media types (text, audio, and video), give context to concepts via hyperlinks, and—ideally—enable personalized learning that is tailored to each student.

Combining these different goals into an interface for learning is a complex task; however, visualization can help (Fig. 2). With Web-based learning systems [11, 17, 18], interactive concept or knowledge map visualizations can promote meaningful learning behaviors [21] while also acting as effective navigation interfaces [29]. These make good progress towards providing easy access to modularized content, in presenting context, and in allowing for some personalization. However, through data, task, and goal analyses conducted with experienced educators, we found that existing systems are yet to create a concept map visualization that is compact, extensible, and supports both linear and non-linear learning plans for different learning styles.

- Michail, Hendrik, James, and Hanspeter are with Harvard Paulson SEAS. E-mail: {schwab, hstrobel, jtompkin, pfister}@seas.harvard.edu.
- Colin is with HarvardX. E-mail: colin\_fredericks@harvard.edu.
- Connor, Dana, Anton, Mayya, and Gary are with the Harvard Institute for Quantitative Social Sciences. E-mail: king@harvard.edu.

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxx

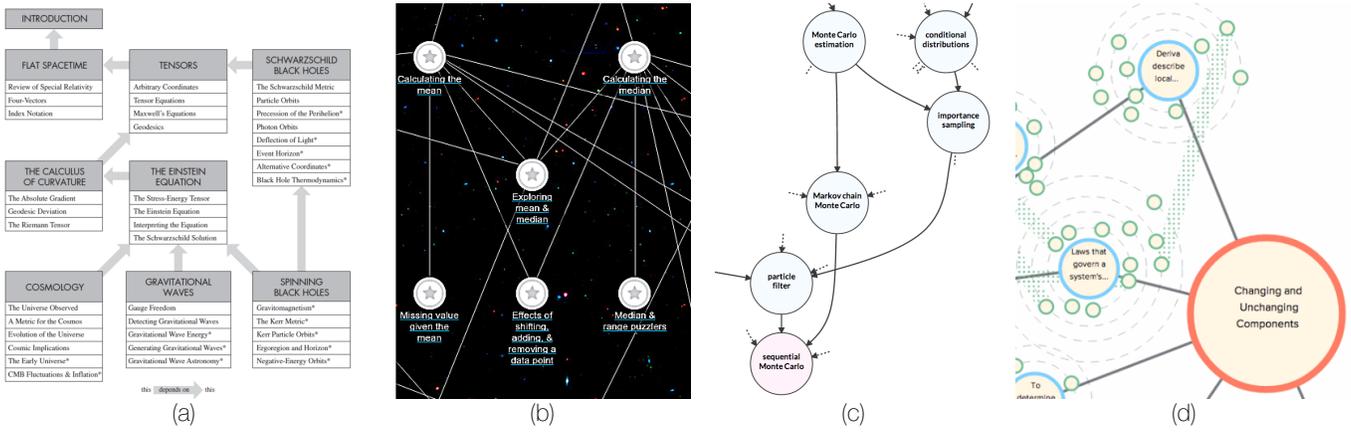


Fig. 2. (a) Non-linear learning dependencies for a book [22], (b) Khan Academy Knowledge Map [17], (c) Metacademy “Deep Belief Network” concept dependencies [11], (d) MIT’s Vinculum [34].

Table 1. Related work feature comparison. *Visualizes Dependencies*: Some systems do not provide concept maps, but those that do show dependencies between concepts. *Visualizes Hierarchies*: MIT Vinculum and booc.io encode concept hierarchies in the visualization. *Guided Learning Plan*: Knewton and booc.io resolve concept dependencies to guide users through courses, but Knewton does not visualize these dependencies. *Different Perspectives*: edX, Knewton, and booc.io provide access to multiple resources for each concept for multiple perspectives. <sup>1</sup>Metacademy collates different sources, though these are not represented on the concept map and must be accessed through a Webpage per concept.

Technique	Examples	Visualizes Dependencies [G3,G4]	Visualizes Hierarchies [G3]	Guided Learning Plan [G4]	Different Perspectives [G2]
Linear syllabus	edX, Coursera, Udacity	No vis.	No vis.	No	Yes
Content recommendation	Knewton [18]	No vis.	No vis.	Yes	Yes
Dependency diagram	Relativity book [22]	Yes	No	No	No
Hierarchical concept map	MIT Vinculum [34]	Yes	Yes	No	No
Traditional concept map	Khan Academy, Metacademy	Yes	No	No	No <sup>1</sup>
Hierarchical circular concept Map	booc.io	Yes	Yes	Yes	Yes

We present *booc.io*, an interactive Web-based learning system that arranges concepts and their dependency tree into hierarchical circular layouts (Fig. 1). Displaying the typically-expansive dependency trees using hierarchical circular layouts makes the design compact and helps provide context. Using zoom interactions allows the hierarchy levels to extend and supports many concepts. Linear presentation of the material, such as is taught in a classroom, is represented as a learning plan that runs “around the clockface” of the hierarchical concept circles. Non-linear learning plans shortcut material and run inside the hierarchy of circles. Personalization is supported by dynamically tailoring learning plans to student progress.

The *booc.io* system was designed from a detailed characterization of domain questions and tasks for educational tools in conjunction with experts. To our knowledge, we present the first combination of concept maps with dynamic, personalized linear and non-linear learning plans. We built a fully-featured system and, with education staff, converted a real classroom course to *booc.io* (Govt. 2001). With this, we conducted a preliminary qualitative evaluation of the system with educators, course creators, and students. Our evaluation of *booc.io* involved twelve students from the real course and eight course and content creators from HarvardX in think aloud studies and structured interviews. We discovered that our concept map visualization was positively received by learners as a useful study tool, which holds promise that our design meets the goals of a Web-supported education system.

## 2 RELATED WORK

Online learning platforms, such as edX, Coursera, or Udacity, help learners to digest material by modularizing content into chunks or

groups, which eases recall by adding structure [1, 9] and allows the most appropriate media type to represent a concept [6]. This improves learner performance and satisfaction [36] especially with active learning and interactive content [26, 4]. Having many media types and hyperlinks to different resources on the Web can aid learning of complex topics by providing multiple perspectives [30, 31], but care must be taken to both provide trust in the material [15] and to prevent learners from becoming *lost in hyperspace* [7].

One way to overcome this confusion is to provide curated linear learning plans: so-called self-guided learning or, with context-aware interfaces, personalized learning [37, 8]. Knewton uses machine learning to recommend assignments from a pool of material based on a learner’s study history [18], but *booc.io* aims to combine them with concept maps—a kind of knowledge graph—to provide context.

Knowledge graphs have long been visualized to express relations [10, 12, 21], and many interactive browsing systems have been built on top of knowledge databases [13]. Some books present visualizations of chapter dependencies as trees to allow learners to plot a plan [27, 22] (Fig. 2a). Online learning platforms often use hierarchical lists to visualize courses. Khan Academy contains a very large tree visualization on a pan and zoom canvas that links many concepts into one map: from “counting with small numbers” to “Euler’s method for solving first-order differential equations” [17] (Fig. 2b). Metacademy allows nodes in the concept map tree to be isolated and for less important edges to disappear for easier assessment of localized context [11] (Fig. 2c). Vinculum (Fig. 2d) integrates hierarchies and dependencies into an atom-inspired representation.

Knowledge can often be arranged hierarchically and, as knowledge dependencies are typically not cyclical, this leads to a tree represen-

tation. The visualization community has long addressed hierarchical tree data (for a broad overview, please see Treevis.net [28] or von Landesberger et al. [35]). This includes developing many sophisticated techniques for space-adaptive visualizations [19]. One example which also shows relations or dependencies is ArcTrees [23], which presents an elongated treemap with external arcs connecting partitions. For large numbers of partitions, the resulting aspect ratio of ArcTrees may be unsuitable, which leads us to discuss circular forms of tree visualizations as a way to generate compact representations.

### Hierarchical Circular Tree Layouts

Bubble or balloon trees [2] use a circle to visually represent each node in a tree, with smaller child circles contained within the parent circle. These circles can be space filling, creating a fractal effect [5]. Circles can be arranged into rings to show more information at once, where child circles of equal weight are placed on the same interior ring [32]. This also allows focusing by placing a child in the center of the largest circle. Bifocal selectors for trees also exist [3]. Circular tree maps can also be combined with node and link diagrams to show information flow [33]. Building upon these works, we investigate zooming and space distortion navigation metaphors for hierarchy exploration, dynamically displaying learning plans through hidden dependency edges, and displaying hierarchical glyphs for personalized progress tracking (Sec. 4.3).

### Concept Maps in Education Research

Visualization of concept or knowledge maps is a key tool across information organization systems [25, 24], and education is no exception [16]. Many works have shown a positive benefit to displaying expert or learner-created maps, and this extends to hypermedia learning as a navigation tool [29]. However, concept maps for an advanced educational course can grow to be very large. Moore tackles the problem of *map shock* when initially viewing complex maps through interaction: by adaptively varying complexity with zooming [21]. With adaptation, the concept map can become the central repository for learning material. Additionally, we enforce compactness by using hierarchical circular trees and adding interactive non-linear learning plans.

We tabulate important facets of common current educational platforms in Table 1. Many systems provide no concept map visualization, which is surprising given the positive effects found in educational research literature, but it gives some indication of the complexity of producing concept maps at scale.

Some systems provide a visual representation of concept dependencies through traditional concept maps. Using hierarchy in the visualization is important to address the aforementioned *map shock*, as well as to provide context. The Khan Academy concept map does not provide this; Metacademy and MIT’s Vinculum [34] attempt to combat this by reducing detail in the edge maps; and Vinculum also provides adaptive hierarchy visualization as per Moore [21]. Using concept dependencies to create a learning plan enables non-linear learning. Traditional concept maps leave this plan creation up to the learner. Knewton’s recommendation system uses internal knowledge of dependencies to propose other concepts for the student to learn, but does not visualize these dependencies to the learner. Concerning multiple perspectives, we judged whether it is easy for students to access additional materials on specific concepts. Typically, specific sites use their own material; Metacademy collates material at a fine grain across specific concepts, though these are not represented on the concept map and must be accessed through a Webpage for each concept.

In summary, there is an opportunity for interactive visualization to better meet learner needs, as no existing system is capable in all areas.

## 3 PROBLEM FORMULATION

Our overall goal is to help meet the needs of online-supported education. To this end, over the course of 18 months, we collaborated with five education experts to develop booc.io. These experts collectively held over 50 years of experience teaching in the classroom at Harvard University and online through HarvardX.

### 3.1 Domain Goals

We began by conducting semi-structured interviews with our experts. Then, we used mind-mapping to find clusters in their comments that defined goals for an ideal system. At a high level, it would act as a repository for material and support students in effectively exploring the material. The system would be able to support educational courses held at a university with residential college students, and with distance learning students. We identified four goals for this ideal system (E1–5 quotes indicate wording used by experts):

**G1: Content Modularization.** Breaking knowledge down to “micro-level units” (E2) and assigning them into 5–9 item groups is essential for memorability [20, 1]. This goal requires expert domain knowledge and should be facilitated by efficient tools. Analogous to a paragraph in writing, each unit should capture one idea or one thought. These thoughts—or *concepts*—could be explained by short videos, excerpts of a book, or scientific articles.

**G2: More is Better.** Spiro and Jehng [30] indicated that students perform better when provided with multiple perspectives on complex topics. While a wide variety of material is available on the Web, students often face the classic problem of being *lost in hyperspace* [7]. To conquer this effect, an ideal learning system should provide means to add content from multiple sources to allow learners to easily access diverse perspectives on each concept. This system “should be a resource for materials—encyclopedia or reference. People using it should see everything” (E1). This should also free content creators from the limit of providing only content which fits linearly into a course or a book, so that including “more information in an organized and accessible way is better than excluding it” (E1).

**G3: Providing Context.** When students learn a concept, they regularly face the question of “why do I need this?” (E3). Knowledge relating concepts to applications can answer this question. While containment of a concept within a group should provide information about similar concepts, dependency information should help answer questions like “where did I come from?” and “where can I go?” Our experts phrased this goal as: “Never lose the big picture—know where you came from and where you can go” (E4).

**G4: Guided Learning Plans.** The system must be able to answer the question of “how can I learn concept X?” (E3). Currently, if a concept builds on concepts taught earlier, the student has to manually create and follow a path through the materials to learn the goal concept as fast as possible. An ideal system should support building an effective and efficient learning plan and show concept dependencies. This plan could be linear, i.e., follow a course structure, or non-linear, i.e., benefit returning students who use the system to catch up with previously-learned material, or revise particular topics. Should the learner be stuck, the system should present appropriate alternative material “at the end of your present knowledge” (E1). This prevents the situation where a single misunderstood concept in a linear presentation “breaks the chain of reasoning” (E1), as alternatives are available. This should be easier than in a linear book or classroom course, where non-linear paths are difficult to access and make it “hard for the motivated learner to learn something technical or difficult” (E1).

### 3.2 User Roles and User Perspectives

During the interviews and meetings, we identified two different user perspectives: the learner and the content creator perspectives.

The *learner* follows a course and learns concepts. She should be helped by an effective visual encoding to form a coherent mental map of the knowledge structure. She will make use of multiple materials within the system to facilitate her understanding of the concepts. Typical user roles are college students, distance students, or motivated individual learners.

The *content creator* creates a course concept hierarchy, defines dependencies, and links learning material to concepts. She should be supported by the system in creating and modifying the hierarchies and dependencies. Ideally, the system will help her re-think the course as a hierarchy of concepts. She should be able to modularize material and assign it to concepts. This user perspective supports the user roles teachers, book writers, or course instructors.

### 3.3 Task Analysis

Based on our analysis of the education literature and existing learning system, our breakdown of these systems into their essential components (Tab. 1), and our domain goals and user perspectives, we developed with our experts a set of tasks that our system will support. The learner perspective will focus on the tasks:

**TLMAP: Visualizing and Navigating a Concept Map.** The hierarchical concept structure should be encoded visually and be navigable with interaction. The representation should be compact and comprehensible to provide context. The learner must be able to depict the linear course structure [G4]. Information about concept dependencies should enrich the visualization [G4]. The user should be able to easily spot visited or already-understood concepts [G1, G3, G4].

**TLMATERIAL: Providing Access to Learning Material.** This should include different content types, and should be scalable to many different materials. To easily explore different perspectives, switching between different material sources should be low effort [G2].

**TLPLANS: Creating and Exploring Learning Plans.** Exploration should be supported by defining a learning goal, reachable via a learning plan and visible on the concept map, that automatically resolves concept dependencies and student progress [G1, G4].

**TLSEARCH: Searching for Concepts.** A crucial and common method for finding concepts should be a search interface. A query will guide a user directly to a concept, but this will be demonstrated within the concept map so that context is not lost [G1, G4].

While the learner perspective is driven by the idea of a content explorer and consumer, the content creator perspective provides mechanisms to build concept structure and add material. The tasks are:

**TCMAP: Creating Concept Maps.** Defining concept hierarchies and dependencies between concepts should be simple. The content creation process should use a similar encoding as the Concept Map, following the WYSIWYG paradigm [G3, G4].

**TCMATERIAL: Adding Course Material.** It must be possible to add many different source types, either in whole or in part. As such, it must be possible to modularize existing course material into segments, with time ranges for external videos or page ranges for book chapters. Following that, material must be assignable to a concept. This process is equivalent to defining edges in a bipartite graph, where a set of concepts is joined many to many to a set of materials [G1, G2].

## 4 DESIGN

The domain goals, user perspectives, and tasks analysis form guidelines to create a visual design for booc.io. First, we introduce our underlying data model (Sec. 4.1) and provide an overview of the major system components (Sec. 4.2). Then, we demonstrate the rationale for the design of the concept map, learning plans, and concept plan components along with a description of user interactions (Sec. 4.3 and 4.4). Finally, we explain the content creator mode (Sec. 4.5).

### 4.1 Data Model

Together with our education experts, we defined terminology to describe our data model and the relations between data entities:

- **Sources** are individual types of material, such as a YouTube video or an ebook.
- **Segments** are excerpts from sources, such as minutes 3:00–7:00 of a YouTube video, or pages 31–37 of an ebook. For additional structure, segments can be formed into **Segment Groups**.
- **Courses** are the top-level collection in booc.io; the organizational structure to support learning.
- **Concepts** are units of knowledge that a course teaches and that a student learns. They can contain any number of segments to describe the concept from different perspectives. Concepts with no segments can be purely structural to form the course hierarchy and to serve as containers for sub-concepts. Typically, the leaves in this tree hierarchy are the “micro-level units”.

The structure relating different data entities is crucial to mapping the data to visual variables (Fig. 3).

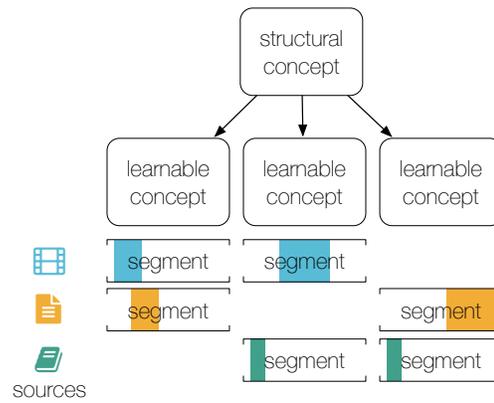


Fig. 3. Course concepts in a hierarchy. Each concept attaches segments from sources, which may be videos, PDFs, or LTIs (Sec. 5).

### 4.2 Overview

Given the hierarchical data model, booc.io must provide a visual encoding for the knowledge tree, plus a way to represent both linear and non-linear learning plans through this tree [TLMAP, TLPLANS]. Which visual metaphor can combine these data characteristics in an easy way without introducing visual clutter? To arrive at our approach, with our five education experts we spent three months theoretically reviewing existing visualization approaches, and evaluating and discussing mock-up conceptualizations against the goals and tasks.

Our design (Fig. 1) is based on a circular layout of the hierarchy as a *concept map*, within which linear and non-linear *learning plans* are presented. To support this map, we also provide a *concept plan* to the left that visually linearizes our circular layout, similar to a file browser, for familiarity. Any interaction or display is mirrored between the map and the plan so that learners can combine their use. Both the learner and content creator perspectives [TCMAP, TCMATERIAL] are supported within these displays, because the concept map and plan both adapt to gain administrative features in content creation mode.

The concept map design is extensible to different situations, including integrating external concepts, visualizing quantitative data, aggregating concepts from across the map into applications, and showing student progress and achievements through iconifying concept circles.

Besides the visualizations, our system includes components for user management and authentication, bulk segment adding and splitting, backups and imports, collaborative content editing, and undo/change history. These features are essential for any learning system and have allowed us to test booc.io with staff and students from a real course (Sec. 6). Because these features are likely less interesting to the visualization community we will refrain from discussing them further.

### 4.3 The Concept Map and Learning Plans

Initially, we iterated over many different designs (Fig. 4) to encode both the hierarchical concept tree [TLMAP] and the linear/non-linear learning plans [TLPLANS]. We discuss how well each design iteration aligns with the system’s tasks and what shortcomings motivated us to design more iterations. With reference to Fig. 4:

- (a) A linear list of concepts with shortcuts is simple, and expands vertically with more concepts, but adding complex non-linear learning plans leads to visual clutter. The expanding size makes it hard to show context. This makes navigating and exploring the concepts difficult [TLMAP].
- (b) Introducing hierarchies with indentation meets the structural data requirement, but it does not solve the expanding space problem and still leads to unwanted clutter. While this is not suitable as a concept map visualization [TLMAP], it is a very familiar structure for browsing trees. We use it as a supplemental *concept plan* for convenience and for fast adoption by new users (Sec. 4.4).

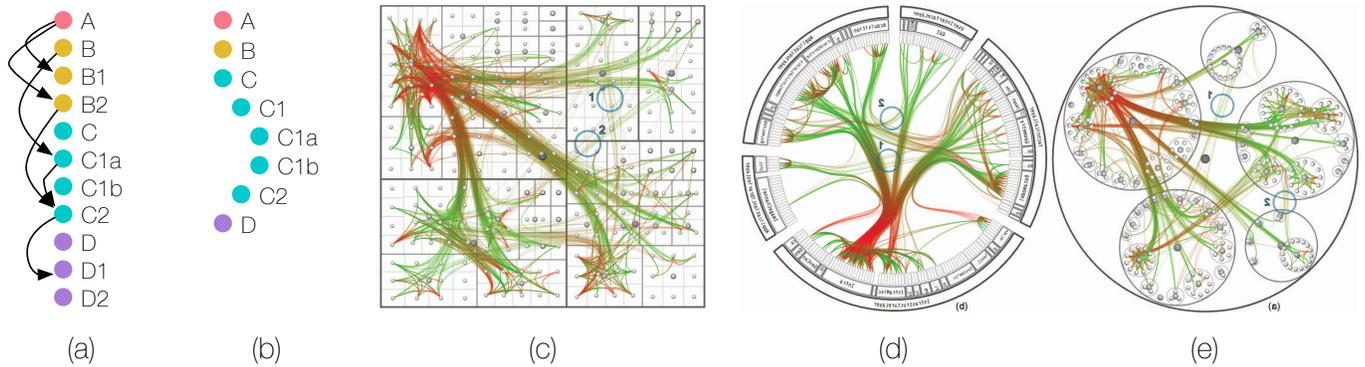


Fig. 4. booc.io design development. (a) Concept list with non-linear shortcuts to represent dependencies. (b) Hierarchical version of (a). (c) Tree map showing a hierarchy with dependencies, which allows a compact visual representation. (d) Radial treemap with dependencies shown internally, which provides better visual segmentation of concepts but still in a compact form, and allows linear paths to run “around the clock”. (e) Zoomable bubble trees make the shape of each subtree more distinctive, increasing contextualization. For our learning plans, dependencies are hidden and are only visible dynamically with interaction. Our initial designs were sketches; we thank Holten [14] for these prototypical versions.

- (c) We considered using space-filling techniques like tree maps, where many modules are packed into a fixed screen real estate for a compact presentation [TLMAP]. However, this imposed new challenges when encoding linear and non-linear paths [TLPLANS], which were hard to interpret when placed on top.
- (d) Circular forms, such as a radial icicle chart, visually separate distinct concept subtrees and are compact [TLMAP]. They also allow us to introduce linear paths through the material naturally with an “around the clock” metaphor. However, linking concepts between different levels of the hierarchy for non-linear learning paths introduces clutter [TLPLANS]. This leaves little space for showing learning materials that should be easy to find [TLMATERIAL].

Inspired by bubble or balloon trees, our final design iteration (e) is a circular layout which encodes hierarchy as sets of child circles. The main differences to a circular tree map are that the circles are not densely packed and that their arrangement within a group is not random. The linear order of child concepts is maintained by using angle as an encoding to place concepts, leading to a “clock-wise” arrangement for the default learning plan [TLPLANS]. This achieves an implicit encoding of linearity without introducing extra lines or clutter. Encoding the hierarchical structure by enclosure (Gestalt principle) and the linear order by angle (visual variable) leaves space to introduce lines to encode non-linear learning plans (Gestalt principle: connectivity) that cut inside the circles. This representation is also compact at the top level to promote context needed for clear navigation of the concepts [TLMAP]. With zooming or space-distortion navigation, it can maintain a 1:1 aspect ratio.

We assign each first-level concept and its contained sub-concepts a specific color hue. For each level of nesting, the color luminance is decreased to indicate and differentiate tree depth. Grey indicates inactive or unnecessary concepts, such as those a student does not need to know for a particular learning plan. Learning plans use the fully saturated version of a color to indicate where the path will end. For example, in Figure 1, the red learning plan leads to the concept *Assessing Models and Estimators*, itself encoded in red. This leads to a clear encoding of the learning plans [TLPLANS], which in turn display the learning material concept by concept [TLMATERIAL].

#### 4.3.1 Interacting with the Concept Map

To navigate the hierarchy of concepts within the fixed screen real estate, we developed two different methods for overview and detail navigation. The first method is based on space distortion (Fig. 5a). A selected sub-concept is enlarged and moved into the center of the concept map, while all other concepts stay in place. This maximizes relevant content while still providing overview. The second method uses a simple click-to-zoom interaction that centers a selected concept and

zooms in, so that the selected sub-concepts fill the screen. Both methods use tooltips to show concept information when labels are no longer visible without zooming or space distortion. While being visually less advanced, this second zooming method was considered to be preferable for a general audience by our expert educator collaborators because the visualization’s stability strengthened the trust in the context.

Finding a concept in booc.io can be accomplished by exploring the map, or by search within the text field in the top left corner [TLSEARCH] (Fig. 1). This allows users to find a specific concept more directly. After completing the search, the chosen concept is brought into focus on the concept map, and the user can explore the local context [TLMAP] and related source segments [TLMATERIAL].

#### 4.3.2 Interacting with Learning Plans

Typically, existing concept maps show all relations or dependencies at once. However, even in simple trees this can quickly lead to visual clutter and makes it hard for learners and content creators to navigate [TLMAP] (Fig. 8d shows the beginnings of dependency confusion during course creation). Instead, we dynamically resolve dependencies into a learning plan when the user hovers over a concept, to directly answer the question “how do I learn x?” [TLPLANS].

Internally, concept dependencies are directed links between two concepts. When selecting a concept, chains of dependencies form a learning plan that leads the user from the beginning of a course to the concept. Knowing the student’s progress, i.e., which concepts she has already seen or understood, allows us to skip or shortcut unneeded concepts. These skipped concepts are desaturated in the concept map.

The hierarchy further lets us simplify the visual presentation of the plan by only showing relevant plan details. For instance, we bypass plan detail at deep tree levels when the user is zoomed out in overview, and instead route the plan through a shallower parent. To inspect the plan in detail, a learner can ‘lock’ or freeze the plan to a goal concept so that it doesn’t change when interacting with other concepts. This allows detailed inspection of the simplified plan (video 00:26).

When no concept is selected, i.e., on first viewing the course, we display the subset of all learning plans that travel from the course root to all immediate child sub-concepts. Typically, these are the major demarcations in a course. This initial display familiarizes the learner with these plans. Each different colored plan is offset slightly so that it can be visually followed from the course start (video 00:45).

In principle, many different learning plan criteria are possible. We currently show the plan that minimizes the number of concepts to learn. However, other strategies might show the linear path defined by the classroom instructor, or consider the length of material within each concept, or even consider popular plans taken by other students.

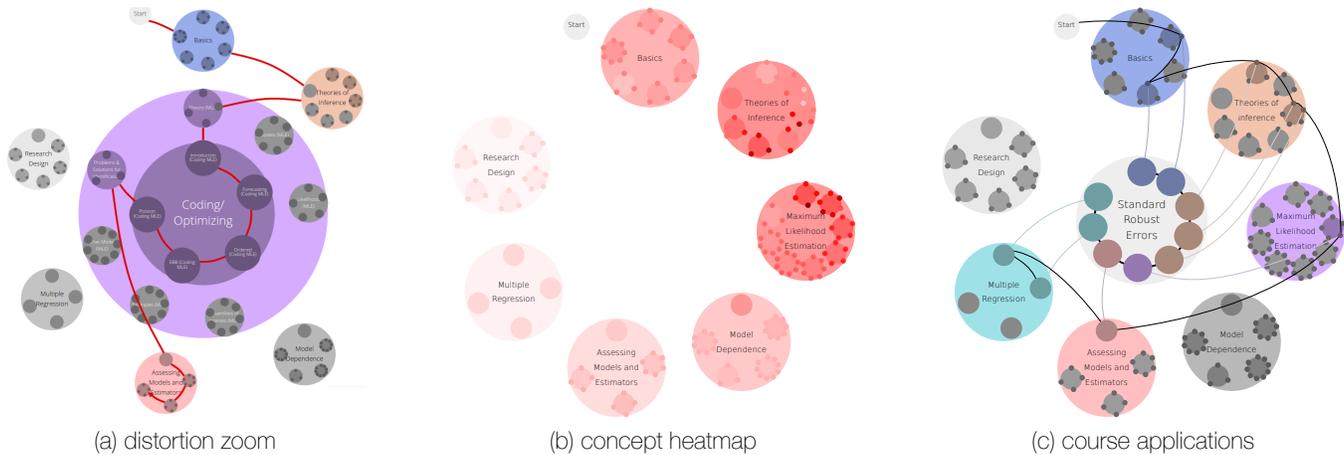


Fig. 5. Concept map design variants. (a) Context and detail by distortion. The concept Coding/Optimizing is selected, and its parent moves from 3rd position around the clock into the center. Then, it moves again into the center from the 4th position around the clock to focus on the desired concept. A red learning plan weaves through the requisite concepts. (b) A heatmap over the concept map displays quantitative course data, such as student activity. (c) Course *applications* fill the central space, and combine concepts taken from across the map using surrogate concepts.

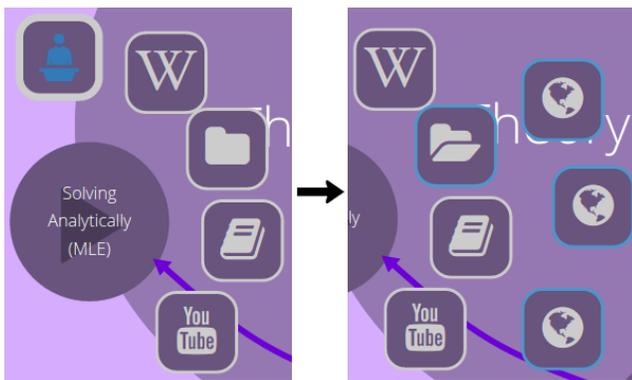


Fig. 6. *Left*: Segments are shown for a selected concept in the concept map, laid out in an arc ring around the concept. Each segment is represented by an icon indicating the type of material (top to bottom: lecture, Wikipedia article, a collapsed segment group, a PDF, and a YouTube video). *Right*: After clicking the group, its contents unfold and are shown highlighted in blue in a second arc.

#### 4.3.3 Interacting with Concept Segments

Each concept contains a set of material segments that appear in an arc around the concept once it is selected (Fig. 6). This avoids overwhelming the user and shows relevant materials when needed [TLMATERIAL]. In the same way that concepts are ordered within their parent, segments are also ordered by angle around the clock, which allows content creators to suggest content to be seen first. It is also possible to visually highlight the first entry to promote material, e.g., material developed specifically for the course. Segment groups are displayed as folder icons, and, once opened, the contained segments form their own second arc around the concept (Fig. 6 in blue outline).

Selecting a segment replaces the concept map by the requested material, e.g., a PDF or video at the start page or time of the specific segment. The learning content is shown in the browser to allow easy context via a quick return to the concept map [TLMAP]. Segment browsing is supported by a learning view that contains common content functionality such as the ability to download content [TLMATERIAL], mark progress, and easily move to the next or previous concepts in the learning plan [TLPLANS].

#### 4.3.4 Progress and Achievements

The concept map design is able to support the easy identification of progress to the user. As the learner consumes content, booc.io auto-

matically tracks their progress and changes the state of concepts from “not seen” to “seen”, which we mark with an eye icon. The user can also mark concepts as “understood”, for which we use a check mark icon. We iconify the parent concept when all sub-concepts are consumed. This forms an easily visible trail of achievements that further helps provide context (Fig. 1) [G3].

#### 4.3.5 Heatmaps, Applications, and External Concepts

**Heatmaps.** The concept map is also designed to support the visualization of qualitative course data as overlays. For instance, it can show heatmaps of aggregated student progress or assessment scores, or discussion forum comment frequency across concepts (Fig. 5b).

**Applications.** While concepts often explain a specific thought at a theoretical level, most application scenarios require a set of independent concepts to be understood. For example, to create a modern web application, a learner needs knowledge about Javascript, Python, web deployment, code versioning, and data wrangling. Applications in booc.io are represented as special concepts containing a collection of *surrogate* concepts that link to many different tree branches (Fig. 5c). We visualize applications by placing the collection of duplicated surrogate concepts in the center of the visualization and draw links to their original locations. Clockwise order is maintained here, too: following the ordering of surrogates within the application concept would allow the learner to understand the application scenario.

**External Concepts.** As the system grows, course instructors will create concepts that have relevance beyond their own course. Instead of creating new instances of, say, a *Maximum Likelihood Estimation* concept, content creators can re-use existing concepts from across other courses in the system. Each course instructor has the option of assigning a priority to order the segments within the concepts of their own course, to promote the most relevant segments.

### 4.4 Concept Plan

The concept plan lies to the left of the concept map, and acts as a search panel, a control panel, and a hierarchical list representation of the course (Fig. 1, left, and Fig. 7a). The concept plan provides familiarity to reduce *map shock*, and gives the interface a “map plus list of directions” often seen in online street maps like Google Maps.

The concept map and concept plan are synchronized to aid map understanding, with each user interaction causing action within both views simultaneously: all displays are mirrored, such as highlighting of selected items, or interactions such as hovering or selecting concepts. For example, searching for a concept name, tag, or material

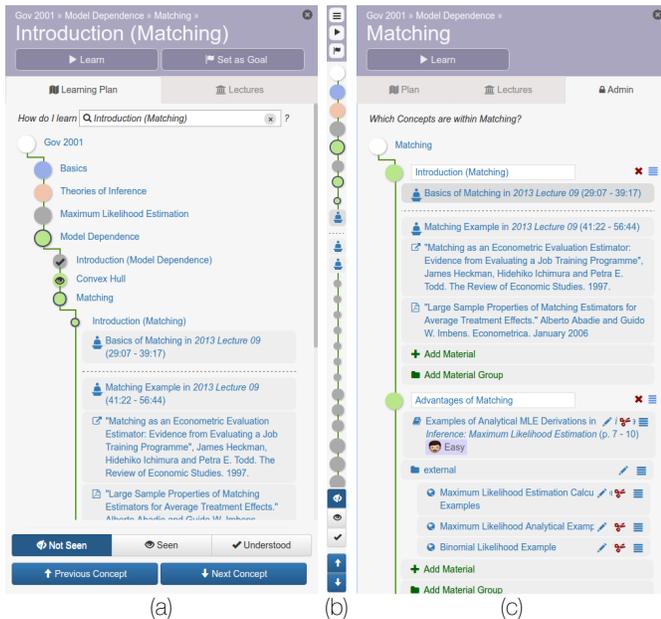


Fig. 7. The concept plan contains a familiar hierarchical list representation of the learning plan. The panel starts expanded (a), but can be minified (b) to give the learner a larger view of the concept map or the current learning material. The concept plan gains edit functionality when viewed from the perspective of a content creator (c). These edits update the concept map immediately. We also see a segment group and its contents, and a tag that has been assigned to the second segment.

type will filter the hierarchical list by collapsing and greying out irrelevant concepts and by hiding irrelevant material; the concept map also greys out to mirror. By default, we maintain course order in the returned search results to provide context [G3], though a more classic search scenario where results are returned by relevance is possible.

If the learner is comfortable with the concept map, the plan can be minified to allow more space for the map or for viewing concept learning material such as videos. The minified version still allows most interactions of the extended version by using a set of symbols (Fig. 7b). In this mode, the linearization of the current learning plan is at its most explicit, represented as a chain of concept circles.

Finally, we provide navigational conveniences on the concept plan: a breadcrumb hierarchy trail below the title, a “hot list” tab that contains the latest lecture or section videos from the classroom for easy access, and the edit tab if the user is a content creator.

#### 4.5 Creating Course Material

For the content creator perspective, all interactions described so far apply, but we add a set of WYSIWYG interactions for content map creation and editing [TCMAP], and for adding and editing course segments [TCMATERIAL]. The edit tab in the concept plan panel shows the course hierarchy with new buttons to add concepts, sub-concepts, and materials and material groups (Fig. 7c). This goes along with the ability to edit the name, color, and presentation order of existing concepts and materials. As the content creator makes edits the concept map is updated immediately (Fig. 8).

To add material segments to concepts [TCMATERIAL], course creators can search and filter through all available sources and their segments, or can create new segments with a dialog customized to each type of material. Course creators can also assign *tags* with optional emoticons to segments to help classify material and provide fast search [TLSEARCH], e.g., to notify students of segment difficulty ( 😬 Hard or 😫 Math ).

### 5 IMPLEMENTATION

booc.io has been built in Javascript using a MEAN Stack: MongoDB, ExpressJS, AngularJS, NodeJS. Learning materials, individual student

progress, and other content is stored on the server using a Mongo database and made accessible through a RESTful API using Express. For the front end we use AngularJS and D3.js. Angular routes users through the application, manages access rights, and keeps the concept plan and other parts of the interface up to date, while D3 is used for the concept map visualization. The arrows are curved using a basis interpolation, and the arrangement of concepts and segments is automatically defined depending on the number of nodes to show.

Users can access booc.io through Google, Facebook, or through Harvard university’s personal authentication (PIN) system. Course staff can add sources of several types, including YouTube videos, PDFs, embedded or external Websites and Wikipedia articles, and lecture slides synchronized with classroom recording videos. The system also supports Learning Tools Interoperability (LTI). LTIs allow content creators to share interactive learning material across platforms, e.g., to embed material from an edX course into booc.io. For this, the system conforms to the OAuth Standard and authenticates users on LTI providers so that any content made available as an LTI will embed without requiring further authentication.

### 6 EVALUATION

To evaluate whether booc.io meets the domain goals and tasks (Secs. 3.1 and 3.3), we took an existing Harvard University course “Govt. 2001—Advanced Quantitative Political Methodology” and asked course teaching fellows to adapt it to booc.io. Then, we conducted two introductory qualitative user studies: 1) a think-aloud study with eight online education experts at HarvardX to gain a perspective on distance learning applications; and 2) structured interviews with twelve student learners who had been enrolled in Govt. 2001.

#### 6.1 Building a Course

The first booc.io requirement is for a hierarchical organization of modular content [G1] to describe concepts, with material (documents, short videos) describing each concept independently. Many existing MOOCs already fit this structure, and their linear organization and material can be imported directly from existing edX courses. This is the ideal situation. For Govt. 2001, the source material was not independent in this way, as it came from longer two-hour lecture videos recorded from the classroom. As such, Govt. 2001 teaching fellows marked timestamps throughout these recordings to segment them into concepts, which took about two days. While machine learning could assist in this task, for instance to suggest similar topics from video transcript analysis, this was beyond the scope of our work.

Next, converting a traditional linear course to booc.io requires the definition of dependencies between content to allow non-linear paths. To provide the most accurate learning plans, this needs to be performed at a fine granularity. This is difficult if the amount of content and the number of potential connections is large; however, the hierarchical organization does reduce this as non-direct parents can be dependencies for child concepts (e.g., the whole of module 2 in a course can be a dependency for a specific concept later on in module 4). Govt. 2001 teaching fellows performed this task, again taking about two days.

Finally, any additional course materials such as supplemental readings and videos were attached to the defined concepts, and the preferential presentation order of content to the learner was set.

#### 6.2 Think-aloud Study

We asked eight HarvardX staff to evaluate booc.io over sessions of 60 minutes. Five of these staff are experienced in creating MOOCs (Massive Open Online Courses) on edX using its Studio interface, which does not use concept map visualizations. Two others are part of HarvardX’s video team, and one was a member of the administrative staff. No staff were familiar with booc.io or Govt. 2001 (i.e., not the five collaborating education experts from Sec. 3), though all were very familiar with the state of online education systems. We asked each individually to tell us their thoughts as they used the system, both from the student and course creator perspectives. At the start of each interview, each participant used the interactive welcome tour. After the

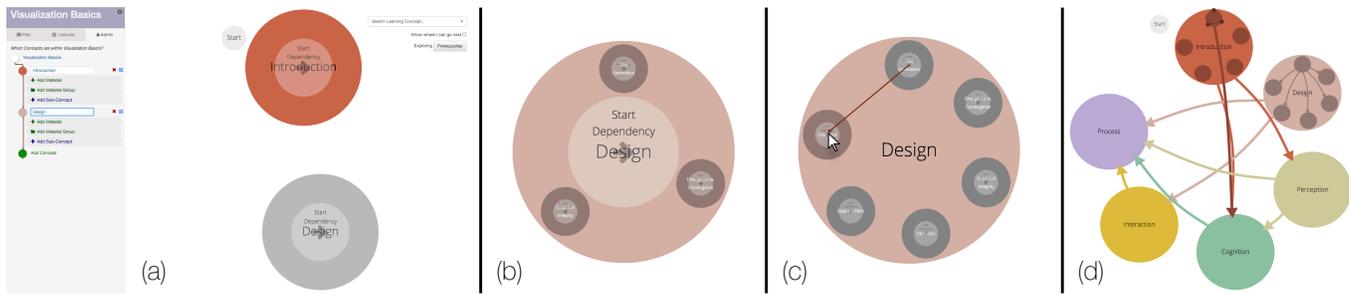


Fig. 8. Course creators begin by adding parent concepts (a) and child concepts (b). Then, dependencies are drawn point to point (c). The zooming interface from learner mode still works as before, making it possible to add dependencies across hierarchies. The final dependency graph may look complex (d), but this is abstracted away via learning plans in the learner perspective.

tour, we let testers explore the system freely before pointing out different features for opinion. During the entire process, we explained almost nothing, and wrote down the participants’ comments (H1–8).

Participants understood the non-linear learning capabilities without being helped, and liked being able to skip some concepts to learn specific concepts (H1: “Oh, I can even skip that to get there!”; H4: “learning plans are cherry picking [in a good way]”, or that H3: “it is good because [the goal system] forces people to have learning objectives”). This strengthens the hypothesis that the visualization effectively shows non-linear learning plans to guide students through the material [TLPLANS]. Participants appreciated the map and understood its context benefits [TLMAP] (H3: “it shows you everything you know and everything you should know—it doesn’t hyper-focus it for you”). One participant initially had trouble with the ‘clockwise’ metaphor (H6: “it seems awkward”), but upon seeing the interactive learning plan was convinced (H6: “Ohhhh, I get it.”). Some participants thought the availability of different perspectives was helpful [TLMATERIAL] (H6: “I like seeing Wikipedia [embedded]”).

For the content creator interface, the ease of content map creation was appreciated [TCMAP] (H8: “it pops up immediately... feels like you’ve accomplished something”) and H5: “this is a heck of a lot faster than what I’m used to [in edX Studio]”). The relationship between defining prerequisites and the dynamic learning plans took more time to grasp [TLPLANS] (H4: “It must draw [the learning plans] based on where the prerequisites are”). Uploading learning materials is simple [TCMATERIAL], but assigning materials to concepts received negative feedback as it used a separate interface. Since then, we moved material assignment directly to the concept plan (Fig. 7).

In general, staff were positive about the system and its possibilities (H2: “Cool concept for organizing stuff.”; H8: “Oh wow, there’s a lot of different ways to use this.”), though there were some reservations about providing choice (H5: “I appreciate the graphical interface, but I really just want someone to push me down the tube of learning”). For booc.io, this tube is following the linear order around the clock.

### 6.3 Govt. 2001 Student Learner Study

We advertised for current or former students of Govt. 2001 to be interviewed about booc.io, rather than random students. This allows participants to evaluate our contributions quickly as there is less barrier to understanding the content, though this may affect their reaction to the system. Twelve participants were randomly selected and paid \$30. Seven were current students of the course, and five had taken the course the year before. The student participants were 23–28, with six females and six males. We interviewed one student at a time and each interview lasted 60 minutes. The students had little to no knowledge about booc.io and were used to using a typical course Webpage containing a hierarchical concept list organized by topic and week with hyperlinks to access course videos and PDFs.

The structured interview proceeded as follows: First, we asked students to find material on a specific concept by using the existing Webpage: “Imagine you are struggling with the likelihood ratio to compare models. How would you use this website to find material on this

topic?” Student habits for finding material on the Webpage included manual search, using a third-party search engine for Govt. 2001 material, or even asking their roommates. Next, we directed students to the booc.io Webpage. Participants followed the standard on-boarding procedure by watching a one-minute welcome video that describes the concept map and learning plan before being shepherded through a two-minute interactive tour of the system within the Govt. 2001 course. After that, participants were given five minutes to familiarize themselves with the system by freely exploring the concept map, non-linear learning plans, and the related materials.

Having learned the basics of the system, we asked participants to find the same learning materials as before in booc.io to see whether the map [TLMAP] and search [TLSEARCH] give students the context needed to achieve typical learning goals. While participants always used the existing Webpage before testing booc.io, which could introduce a potential learning effect of ‘knowing where the content is’, in practice the booc.io content is both chunked under different titles and under a different organizational structure. After completing the task, we prompted students to explore other booc.io features for about ten minutes. Finally, we interviewed students with set questions about their impressions and experience using booc.io (S1–12).

In post-interview analysis, we transcribed the interview videos and classified all comments by students into positive, neutral, and negative categories for each specific feature to try to assess how well the system performed for each goal and task (Secs. 3).

### Results

The multiple levels of hierarchy and context from the interactive concept map [TLMAP], many types of learning material [TLMATERIAL], personalized learning plans [TLPLANS], and search [TLSEARCH] were all well received.

[TLMAP]: All twelve students said they liked the concept map, and many mentioned that it provided context, overview, or even seeing the “overarching idea” (S12) or “a sense of the course architecture” (S11). It also clarifies the context: “Very useful to see how concepts relate to each other. That helps not to confuse things that are taught in succession” (S8), confirming the map’s role in providing context [TLMAP]. Some *map shock* still existed: “so many bubbles” (S12); in response we removed defaulting to showing future ‘where can I go from here?’ edges. Two students disliked the concept map *for navigation*, even if they liked it for providing context. These students specifically self-identified as ‘non-visual learners’: “I don’t love the map, visually. It’s not how I learn. But it’s useful to see the bubbles to see the hierarchy” (S3). They preferred text as a navigation structure, for which the concept plan was well received as it duplicates all map functionality in a more familiar hierarchical list.

[TLPLANS]: All twelve students stated that seeing the fastest, non-linear plan to a given concept as well as relationships between concepts is useful. They understood the role of the learning plans quickly and appreciated their value (“It was pretty clear that in order to understand negative binomial model, you need to understand these things that are connected” (S1), “The connected circles idea is great because

it shows me exactly what I need to know” (S2)). As an organizational structure, one student said “I wished I had this in general in life because I have often wished I had something that could tell me how to do certain techniques by knowing what I need to know” (S9).

[TLMATERIAL]: Being able to chose multiple materials for different perspectives was also well received with all twelve students’ endorsements, with some saying that they would “otherwise have to find these resources on their own” (S1), and that “the teachers curating the contents assures trust in the material” (S9). However, some *map shock* existed here as well: “that’s a lot of options”. Careful use of segment groups can help to not show too many materials at once.

[TLSEARCH]: The idea of a search feature was very well received, although two students remarked that the search was not powerful enough. One student said she would like to be able to “filter material or concepts by *application*, *theory* and other labels” (S12), and this has since been added to the system.

When asked to find material on the likelihood ratio using the existing course Webpage, current students of the course took a median time of 1:02 min., measured as the elapsed time between posing the question and finding the material. Students mostly used their memory of the lectures to recall how far back the content was chronologically, then keyword searched within the relevant week’s PDF slides. Of the five students who took the course previously, four gave up and used a Web search. When the twelve students were asked to perform the same task on booc.io, six of them used the search [TLSEARCH], and six used the visualization to navigate within the concept hierarchy [TLMAP]. The median times spent were 25 sec. and 30 sec. respectively, which validates the importance of efficiency in these tasks. All five former students of the course found the relevant materials using booc.io, while only one of them managed to do so on the existing Webpage. This shows that less material familiarity is needed with booc.io than with the existing Webpage.

During the user study we became aware of several improvements that would benefit the system. To help support in-person courses, participants suggested that it should be possible to highlight or filter map concepts that had already been covered in class. Students also requested much more powerful search tools, for instance, to be able to look within segments for specific phrases used in class. One aspect not addressed by our system but requested by participants was for the “understood” progress tracking to be tied to formal assessment. The idea of marking a concept as “understood” was often seen as too coarse. In principle, our design could be extended by using a unique icon to identify assessed concepts and material. While an assessment might consist of a grade, students also suggested two alternatives: a Likert-like scale for self-assessment of their comprehension, or self-assessment based on short multiple choice tests.

## 7 DISCUSSION AND FUTURE WORK

booc.io helps learners explore a hierarchy of concepts and different learning plans through them. An essential prerequisite for this is the proper modularization of course content [G1]. Creating concepts as self-contained units that can be consumed independently is a challenge for content creators. It requires a shift in perspective from creating a 60–90 minute long linear lecture to creating hierarchical “micro-level units”. While we provide tools to segment existing material to help bridge the gap, we also see our system as a tool to motivate and support this shift. This sentiment has been echoed by our HarvardX collaborators: providing a compelling learning interface helps content creators to think about the modularization of their own courses and to motivate them to create new material. In distance learning, this trend is exemplified in MOOC presentation styles. In this context booc.io manages to meet more of our identified goals than existing systems.

In principle, booc.io provides a platform for democratizing course material. While it is common for teachers to create their own course materials for a specific concept, it is also common for students to look for and benefit from secondary resources that explain the concepts from different perspectives [G2]. We imagine that a future version of booc.io will contain a large pool of concepts contributed by various experts. Course instructors will be able to sample a useful set of con-

cepts to create a course, add their own segments, and recommend to students curated lists of alternative explanations. We also imagine students will be recommending content for inclusion. This again requires a shift in perspective: to decouple the creation of segments from the creation of concepts, and to be willing to share this material (though this is less problematic in academia). For instance, a visualization course teacher could use and extend concepts on *Bertin’s Marks and Channels* to help build a set of tried and true explanations.

In this regard, as more segments and concepts are added, any hierarchical representation has scalability trade-offs, and booc.io is no exception. The number of sub-concepts that can be arranged around the circle of a super-concept is limited to about nine such that sub-concept circles do not become too small. This number roughly corresponds to the number of items that humans can hold in short-term memory ( $7 \pm 2$ ) [20, 1]. Practically, in our real course, our Govt. 2001 course creators rarely used this many sub-concepts, and it was always possible for them to re-group some into another tree branch.

There is an opportunity to integrate social features into booc.io, including social networking elements like concept and material discussion, or *like* functions for segments that explain concepts very well. Tracking students and their interactions within the system would be useful input to machine learning models to predict favored learning plans or favorable segments. The segments a learner digests before marking a concept as understood could be an indicator for a good explanation, which the system would try to identify automatically and promote to other learners—a ‘wisdom of the crowd’ approach to building trust in material. A predictive model could be trained to recommend certain types of segments for a given history of learning [18].

While booc.io will help students learn specific concepts, their feedback suggests that we could better integrate the system for in-class use. Having lecturers refer to the “big picture” concept map in class can immediately motivate topics, particularly applications. As one participant said: “Oh, the big picture is shown—that never happens in class!” (S6). The learner and concept creator perspectives could also be augmented by a *teacher* perspective: this could provide a concept map with additional highlighting and digital annotations so that attention could be drawn to sets of concepts of current importance.

The visualizations we created target an audience with many different backgrounds, where most users are not used to seeing state-of-the-art visualizations. The broadness of the user base motivates the use of simple visual metaphors, e.g., by choosing context + detail zooming over space deformations (Sec. 4.3.1). As one student put it, “I’m not a tech geek but I still like this” (S11).

With only one full course built with booc.io, we must be careful in generalizing suitability. However, our expectation is that the system is flexible to different courses, such as those imported from edX and augmented with dependency information. Anecdotally, professors from across different disciplines have wished to use booc.io on first impression (e.g., quantitative biology, astrophysics).

Finally, our current evaluation is only introductory and our participant results should not be unduly extrapolated. As an education system with broad scope, we require an in-depth study of learner and content creator use over many courses. This is beyond the scope of this work, but we intent to expand the evaluation to encompass longitudinal study feedback and a broader set of courses.

## 8 CONCLUSIONS

We have presented booc.io, an online education system that combines hierarchical concept maps and dynamic non-linear learning paths. Through interviews with education experts we identified several important domain goals and learning tasks. Our design addresses each of these tasks by visually encoding the hierarchy of concepts as well as linear and non-linear learning plans. The system encourages including lots of material on individual concepts to facilitate multiple perspectives. In a preliminary evaluation, students from a Harvard course provided positive feedback and many suggestions for future work. booc.io brings together and extends different aspects of existing learning systems to meet the needs of modern online education, and it holds promise as a useful tool for both learners and course creators.

## REFERENCES

- [1] A. D. Baddeley and G. Hitch. Working memory. *Science*, 255:556–559, 1992.
- [2] R. Boardman. Bubble Trees – the visualization of hierarchical information structures. In *CHI'00: Extended abstracts of the SIGCHI conference on Human Factors in Computing Systems*, pages 315–316. ACM Press, 2000.
- [3] R. A. Cava, P. R. G. Luzzardi, and C. M. D. S. Freitas. The bifocal tree: A technique for the visualization of hierarchical information structures. In *IHC'02: Proceedings of the Workshop on Human factors in Computer Systems*, 2002.
- [4] M. T. H. Chia and R. Wylie. The ICAP framework: Linking cognitive engagement to active learning outcomes. *Educational Psychologist*, 49:219–243, 2014.
- [5] G. Clifton-Bligh. Methods and devices for mapping data files, 2004. US Patent 6775659 B2.
- [6] D. Cunningham, T. Duffy, and R. A. Knuth. The textbook of the future. In C. McKnight, A. Dillon, and J. Richardson, editors, *Hypertext: A Psychological Perspective*, chapter 3, pages 19–50. Ellis Horwood, New York, 1993.
- [7] D. M. Edwards and L. Hardman. Lost in hyperspace: Cognitive mapping and navigation in a hypertext environment. In R. McAleese, editor, *Hypertext: theory into practice*, chapter Lost in Hyperspace: Cognitive Mapping and Navigation in a Hypertext Environment, pages 90–105. Intellect Books, Exeter, UK, UK, 1999.
- [8] J. Ferreira. Methods, media, and systems for computer-based learning, Feb. 19 2009. US Patent App. 12/192,112.
- [9] H. Freyhoff, H. Gruber, and A. Ziegler. Expertise and hierarchical knowledge representation in chess. *Psychological Research*, 54:32–37, 1992.
- [10] B. R. Gaines and M. L. G. Shaw. Concept maps as hypermedia components. *Int. J. Hum.-Comput. Stud.*, 43(3):323–361, Sept. 1995.
- [11] R. Grosse and C. Reed. Metacademy Deep Belief Networks Concept Map. [https://www.metacademy.org/graphs/concepts/deep\\_belief\\_networks#mode=explore](https://www.metacademy.org/graphs/concepts/deep_belief_networks#mode=explore). Last accessed 28th March 2016.
- [12] I. Herman, G. Melancon, and M. S. Marshall. Graph visualization and navigation in information visualization: A survey. *IEEE Transactions on Visualization and Computer Graphics*, 6(1):24–43, Jan 2000.
- [13] C. Hirsch, J. Hosking, and J. Grundy. Interactive visualization tools for exploring the semantic graph of large knowledge spaces. In *Workshop on Visual Interfaces to the Social and the Semantic Web (VISSW2009)*, volume 443, 2009.
- [14] D. Holten. Hierarchical edge bundles: Visualization of adjacency relations in hierarchical data. *Visualization and Computer Graphics, IEEE Transactions on*, 12(5):741–748, 2006.
- [15] M. K. Iding, M. E. Crosby, B. Auernheimer, and E. Barbara Klemm. Web site credibility: Why do people believe what they believe? *Instructional Science*, 37(1):43–63, 2008.
- [16] O. O. A. John C. Nesbit. Learning with concept and knowledge maps: A meta-analysis. *Review of Educational Research*, 76(3):413–448, 2006.
- [17] Khan Academy Knowledge Map. <https://www.khanacademy.org/exercisedashboard>. Last accessed 28th March 2016.
- [18] Knewton. <https://www.knewton.com/>. Last accessed 28th March 2016.
- [19] J. Lamping, R. Rao, and P. Pirolli. A focus+context technique based on hyperbolic geometry for visualizing large hierarchies. In I. R. Katz, R. Mack, L. Marks, M. B. Rosson, and J. Nielsen, editors, *CHI'95: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 401–408. ACM Press/Addison-Wesley Publishing Co., 1995.
- [20] G. A. Miller. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63:81–97, 1956.
- [21] J. P. Moore. *Promoting Conceptual Understanding via Adaptive Concept Maps*. PhD thesis, Blacksburg, VA, USA, 2013. AAI3585800.
- [22] T. A. Moore. *A General Relativity Workbook*. University Science Books, 2012.
- [23] P. Neumann, S. Schlechtweg, and M. T. Carpendale. Arctrees: Visualizing relations in hierarchical data. In *Proceedings of Eurographics, IEEE VGTC Symposium on Visualization (EuroVis2005)*, pages 53–60. Eurographics, 2005.
- [24] J. Novak. *Learning, Creating, and Using Knowledge: Concept Maps as Facilitative Tools in Schools and Corporations*. Routledge, 2010.
- [25] A. M. O'Donnell, D. F. Dansereau, and R. H. Hall. Knowledge maps as scaffolds for cognitive processing. *Educational Psychology Review*, 14(1):71–86, 2002.
- [26] M. Prince. Does active learning work? A review of the research. *Journal of Engineering Education*, 93:223231, 2004.
- [27] S. Prince. *Computer Vision: Models Learning and Inference*. Cambridge University Press, 2012.
- [28] H. J. Schulz. Treevis.net: A tree visualization reference. *IEEE Computer Graphics and Applications*, 31(6):11–15, Nov 2011.
- [29] R.-S. Shaw. A study of learning performance of e-learning materials design with knowledge maps. *Computers & Education*, 54:253264, 2010.
- [30] R. J. Spiro and J.-C. Jehng. Cognitive flexibility and hypertext: theory and technology for the nonlinear and multidimensional traversal of complex subject matter. *Cognition, Education, and Multimedia: Exploring Ideas in High Technology*, 1990.
- [31] K. Swan. Learning effectiveness online: What the research tells us. *Elements of Quality Online Education, Practice and Direction*, pages 13–45, 2003.
- [32] S. T. Teoh and K.-L. Ma. RINGS: A technique for visualizing large hierarchies. In M. T. Goodrich and S. G. Kobourov, editors, *GD'02: Proceedings of the International Symposium on Graph Drawing*, Lecture Notes in Computer Science, pages 51–73. Springer, 2002.
- [33] F. B. Viégas, M. Wattenberg, J. Hebert, G. Borggaard, A. Cichowlas, J. Feinberg, J. Orwant, and C. R. Wren. Google+ Ripples: A native visualization of information flow. In *WWW'13: Proceedings of the International World Wide Web Conference*, pages 1389–1398. ACM Press, 2013.
- [34] MIT Teaching and Learning Lab: Vinculum. <http://vinculum.mit.edu/>. Last accessed 28th March 2016.
- [35] T. von Landesberger, A. Kuijper, T. Schreck, J. Kohlhammer, J. van Wijk, J.-D. Fekete, and D. Fellner. Visual analysis of large graphs: State-of-the-art and future research challenges. *Computer Graphics Forum*, 30(6):1719–1749, 2011.
- [36] D. Zhang, L. Zhou, R. O. Briggs, and J. F. Nunamaker. Instructional video in e-learning: Assessing the impact of interactive video on learning effectiveness. *Information & Management* 43, 43:1527, 2006.
- [37] J. Zumbach. Cognitive overhead in hypertext learning reexamined: Overcoming the myths. *Journal of Educational Multimedia and Hypermedia*, 15:411–432, 2006.