Chapter 12
Scripting and Orchestration of Learning Across Contexts: A Role for Intelligent Agents and Data Mining

Mike Tissenbaum and James D. Slotta

Abstract This chapter describes a 12-week physics curriculum that engaged students as a knowledge community across contexts: in their classroom, home, neighborhoods, and in the smart classroom. In order to support the curriculum intervention, we developed two complementary technology environments that are built on SAIL Smart Space (S3) – an open-source technology framework. Using a design-based research methodology, we instantiated an orchestrational framework that included the use of social tagging and metadata. Additionally, we devised intelligent agents to support the enactment of our collaborative inquiry scripts. We identify three important structural dimensions for which intelligent agents can play a key role in the orchestration of such curricula: Content Agents, Activity Structure Agents, and Grouping Agents. We conclude with an evaluation of these agents in support of our curriculum designs and propose a set of design principles for the role of intelligent agents and data mining in supporting cross-context learning.

Introduction: New Opportunities for Learning Across Contexts

In recent years, we have witnessed a change in the ways in which students are engaging with the world around them. Over two thirds of Americans now have Internet access at home, and the majority of teens are now actively engaged in the creation of online content (Bull et al. 2008). Outside of school, students are increasingly driving their own learning, by finding relevant resources or connecting to online interest groups, using Internet or cellular network technologies to mediate their interactions (Sefton-Green 2004). These new practices are familiar to students, who have grown up with a "Web 2.0" landscape, where users are the active creators, commenters, and classifiers of the products and processes with which

---

M. Tissenbaum (✉) • J.D. Slotta
University of Toronto, Toronto, ON, Canada
e-mail: miketissenbaum@gmail.com

L.-H. Wong et al. (eds.), *Seamless Learning in the Age of Mobile Connectivity*,
DOI 10.1007/978-981-287-113-8_12

they interact, including the construction and organization of knowledge (Dohn 2009). Such user-created content can take on many different forms, from collections of user-contributed artifacts (e.g., Flickr, YouTube), to community-generated social spaces (e.g., Facebook, ResearchGate), to collaboratively generated and edited evolving content (Wikipedia), to news feeds or other socially filtered resource streams (e.g., Reddit). Even games and leisure spaces are now deeply infused with a social component (e.g., World of Warcraft, Fantasy Sports). There is evidence that such "user-contributed content" (Vickery and Wunsch-Vincent 2007) promotes deeper engagement with the content and the learning community, because users see themselves as participating in the community's progress (Tedjamulia et al. 2005) and because of the awareness of "having an audience" (Wheeler et al. 2008). Despite the explosive growth of these practices in many domains, and their increasing importance to everyday life in the twenty-first-century knowledge society (Zuboff and Maxim 2004), schools have generally failed to adapt them into regular curricular designs (Buckingham 2007).

The growth of user-contributed content is paralleled by the rapid advancement of technologies that support users in connecting and participating in these communities. Common to many of these technologies is the use of metadata (i.e., data about data; Wiley 2000), which can be both user generated or system generated. User-generated metadata often takes the form of tagging, in which participants assign keywords to objects within the system (e.g., photos, videos, narratives, or other users). System-generated metadata can be generated automatically to capture complex underlying information about both the users of the system (e.g., assigned roles, group memberships, times logged into the system) and the products of their interactions (e.g., created artifacts, votes cast, pages visited). Connecting this metadata to semantic ontologies (e.g., to well-defined categories such as "tags," "groups," or "roles") can provide avenues for connecting seemingly disparate pieces of information within a community's knowledge base (Anderson and Whitelock 2004). This semantic metadata can also be leveraged to coordinate access to materials and activities, group assignments, and other logical functions of the system (Simon et al. 2004; Zhao and Okamoto 2011).

The use of metadata to make connections between individual students and the products of the larger class community becomes particularly powerful when researchers wish to extend learning beyond the traditional classroom walls. Social and semantic metadata can create a "chain" that connects student learning across formal and informal learning contexts (in class, at home, in the field, or in their neighborhoods) and across diverse time scales versus traditional single class periods (Milrad et al. 2013). Sometimes referred to as "seamless learning" (Chan et al. 2006), this approach can empower learners to engage with their learning community's knowledge base, whenever and wherever they are situated (Wong and Looi 2011). With more than two thirds of young adults now owning a web-enabled smartphone in the United States (Pew 2012[1]), there is a growing technological capability for students to engage with their learning community "on the go." Such

---

[1] http://www.pewinternet.org/Reports/2012/Smartphone-Update-Sept-2012.aspx

capabilities are not in themselves sufficient to ensure effective learning designs, but fortunately they also provide new opportunities for research of such learning. Because research has typically focused on either formal or informal learning and not on the synergistic connection across contexts and environments (Looi et al. 2010), a central goal of seamless learning is to develop curricular interventions where students access and contribute to community knowledge across diverse learning contexts and scales of time (Wong and Looi 2011). This challenge entails the design of user interfaces, including the representation of community knowledge, and how information is used and contributed meaningfully by students across distinct contexts.

Scripting and Orchestration

Curricular designs that include student-contributed content and learning across contexts are likely to be more complex and dynamic than in previous generations of computer-supported learning (Slotta 2010). Designs must now include the configuration (and possibly the dynamic reconfiguration, based on emergent metadata) of student groups and activities, the technologies used, and critical roles for the teacher. Even in traditional classroom settings, when left to their own devices, students often struggle to choose the most appropriate strategies, understand the goals, or the nature of the task (O'Donnell and Dansereau 1992). Thus, in designing curricula that span multiple intertwined learning contexts, we must carefully configure the structure of the interactions, roles, goals, and interaction patterns in the form of pedagogical "scripts" (Kaplan and Dillenbourg 2010). These scripts have a macroscopic aspect that describes the overall curriculum and timing of individual activities (e.g., a field trip, or a homework task) and a microscopic aspect that specifies individual activities at the fine-grained detail of specific materials, tools, and learning goals (Tissenbaum and Slotta 2012). The design of both the macro- and micro-scripts must address the content of the learning domain, including the specific learning goals for students.

The enactment of such carefully designed scripts is typically scaffolded with computer-based learning environments such as WISE (Slotta and Linn 2009) or with scientific experimentation environments such as Vlab (Tsovaltzi et al. 2008). With the growth of mobile technologies, new technology environments have been developed to support student observations in museums (Kuhn et al. 2012), university campuses (Kohen-Vacs et al. 2011), or environmental field trips (Zimmerman and Slotta 2003). A fundamental challenge to seamless, cross-context learning will be the integration of such learning environments, allowing students to experience a productive, engaging "macro-script" that includes distinct micro-scripts within each context.

There is a parallel challenge of coordinating students and teachers during the enactment of such scripts. This process of supporting the execution of these scripts, both in real time and across longer scales of time, is often referred to as orchestration

(Dillenbourg et al. 2009). The orchestration of curricular scripts needs to be flexible enough to allow for the emergence of new ideas, themes, and avenues for investigation (Slotta 2010). As these scripts become more complex, the information processing needs of both teachers and students increase significantly, requiring designs to consider the "orchestrational load" of participants (Dillenbourg et al. 2011). In managing the orchestrational load of the classroom, designs need to take into account the various actors' informational and regulatory needs and distribute this load among the participants, materials, and technologies present within the learning environment (Sharples 2013). Students need to make sense of their place within the script, their role within the class, and how to access relevant materials within an evolving knowledge base. Similarly, teachers must be aware of what is happening on individual, small group, and whole class levels; the timing and progress of activities; the state of knowledge within the class; and potential points of intervention within the script. To respond to such challenges, technological supports must support the flow of materials, the scaffolding of activities, and the real-time processing of user interactions to inform student (self-regulated)- and teacher-mediated orchestration. Cuendet and Dillenbourg (2013) suggest that the use of distributed interfaces, spreading orchestrational information and regulatory process across multiple interfaces, can be a successful strategy for reducing orchestrational load. This approach becomes especially powerful in smart classroom designs in which the multiple modalities of the room (e.g., tablets, screens, interactive tables and walls, paper artifacts) all disappear into a single unified classroom ecosystem (Cuendet and Dillenbourg 2013).

Intelligent Agents for Scripting and Orchestration

One approach to such orchestration is seen in the application of "intelligent software agents" – small, active software elements that can respond to current context, or past actions of participants, performing real-time data-mining operations and operating on semantic metadata (Brusilovsky 2001). For example, the assignment of students to groups and the assignment of materials to groups can be informed dynamically by processing the metadata of what materials students have worked on previously or their location within the physical environment (Tissenbaum and Slotta 2013). Intelligent agents hold particular promise in support of inquiry learning, in part because they allow orchestration of scripts that are deliberately ill determined at the outset of orchestration (i.e., scripts where it is not known, a priori, what outcomes or conditions will emerge from the products of student interactions). The use of intelligent agents allows for such open-ended designs that allow the script to evolve in relation to student interactions (Slotta 2010).

These new forms of evolving curricular scripts are well suited to the design and enactment of activities where students contribute to and make use of the growing knowledge base in learning activities across multiple contexts. To the extent that any learning activities depend on student-contributed materials, it is not actually possible to know in advance the complete content or structure of such activities. Metadata, such as student-generated tags or votes, will emerge as a result of the enactment, and activity sequences may be scripted such that they depend on those emergent features. We identify three important dimensions of structure, for which intelligent agents can play a key role in the orchestration of such curriculum:

1. Content Agents
   This refers to the use of intelligent agents for managing, building, and retrieving content. What is the current domain of a student's inquiry, and what learning context, group, or tool are they learning? By understanding the content that students are, or have been, working on, intelligent agents can update students on changes to that content or connect it to other artifacts for knowledge work. Agents also have the opportunity to inject materials into the script, e.g., by populating a student's "drawer" (within a particular learning environment) with all content materials that are tagged by students – even those appearing in real time.
2. Activity Sequencing Agents
   As student- and system-generated semantic metadata emerge, data-mining agents can connect users with materials, as described above, but can also make assignments to learning activities or conditions. Sequencing agents can process a student's interactions while also monitoring global (i.e., community level) metadata, to determine the next activity, tool, or location for the student. In this way, the script does not have to be identical for all students and can be seen more as a map of activities that students can traverse in many pathways. Sequencing agents help determine what parts of the map may be accessible, in accordance with emergent metadata and scripting logic.
3. Grouping Agents
   The ability to know the history of student interactions, both individually and as part of the larger community, allows for the design of intelligent agents that can dynamically group or sort students according to specific pedagogical logic. This has particular significance in managing the orchestrational load for teachers, by helping track and manage which students have worked with whom, what materials students have covered in past activities, or any groups (e.g., tasks or expertise groups) to which they have previously been assigned. Intelligent agents can group students with peers according to metadata that is emerging in real time – which would be practically impossible for any human to do in real time.

New Pedagogical Models for Collective Inquiry

The new technology affordances and pedagogical constructs described above present challenges for teachers and researchers to design pedagogical applications that include user-contributed content and cross-context learning. Educators require clear models of how to engage students in such learning, connecting their personal activities meaningfully into a larger social construct and supporting activities across long spans of time and various contexts. One approach to the integration of Web 2.0 tools and practices is that of knowledge communities, where students are asked to see themselves as a collective learning unit, with a high level of responsibility for defining their own learning goals and activities (Brown and Campione 1996; Scardamalia and Bereiter 1994; Bielaczyc and Collins 1999). In a knowledge community, students contribute content to a central "knowledge base" where it is accessible to all peers in the community as a resource for subsequent inquiry activities (Slotta and Najafi 2010). Ideas can also be refined or improved by members or synthesized into higher-order learning objects (Scardamalia and Bereiter 2006). Despite its clear relevance to the needs of twenty-first-century learning, the knowledge community approach has not been widely adopted by teachers or researchers, due in part to the high demand that it places on teachers. This is particularly true in domains with substantive demands for content learning, such as secondary math and science, where teachers do not feel that they have the luxury of encouraging their students to work as a knowledge community and define their own learning objectives (Slotta and Peters 2008).

To make the knowledge community approach more accessible to teachers, Slotta and his colleagues have developed the Knowledge Community and Inquiry (KCI) model (Fig. 12.1), which specifies a set of design principles for a knowledge community approach for secondary science (Slotta and Najafi 2013). In KCI, students work collectively, contributing, tagging, and improving content in a shared knowledge base that serves as a resource for subsequent inquiry. Inquiry activities are carefully designed so that they engage students with targeted content and provide assessable outcomes, allowing students some level of freedom and flexibility but ensuring progress on the relevant learning goals. KCI curriculum requires a substantive epistemic shift away from didactic presentation of content (where students work

Fig. 12.1 Knowledge community and inquiry (KCI) model

largely under the guise of individual learners) and toward a collective understanding of progress and activity. KCI guides the design of complex inquiry activities that span multiple student configurations (individual, small group, and whole class) and contexts (in class, at home, in the field). Within a KCI curriculum, which can span weeks or months, students explore and develop ideas using technology-enhanced materials, tools, and interactive simulations. These activities are carefully scripted in order to address specific learning goals; however, the script itself must be flexible enough to allow for the emergence of new ideas and community voice. This may include the development of technologies that allow the teacher to easily author new activities, scaffolds, or prompts in response to these emergent factors. As part of the development of KCI, we have established a set of design principles that guide the creation of individual, cooperative, collective, and collaborative scripts and activities and how these are scaffolded.

New Technology Supports for Collective Inquiry

In order to support the curriculum where students are engaged collectively in sustained inquiry, researchers have developed technology environments that scaffold student learning and support the growth and procedure of the knowledge community. Scardamalia and Bereiter (2006) have argued that such environments must be aligned with the underlying epistemic goals of the approach. They developed an environment called Knowledge Forum, largely because existing technologies could not support the types of interactions demanded by their theoretical principles of knowledge building. Thus, any technology environments employed within such a knowledge community approach are more than just tools or workspaces, but, rather, they serve as an integral part of the community's distributed intelligence and are intrinsic to the community's notions of learning and intelligence (Pea 2004).

In order to successfully enact the kinds of complex designs required for KCI, we needed a flexible and adaptive infrastructure that could support the design and orchestration of collaborative activities that include spatial, social, and semantic dependencies. To this end, we have developed SAIL Smart Space (S3), an open-source framework that coordinates complex pedagogical sequences, including dynamic sorting and grouping of students and the delivery of materials based on emergent semantic connections. S3 has been developed to allow the physical space of classrooms or other learning environments to play a meaningful role within the learning design – either through locational mapping of pedagogical elements (e.g., where different locations are scripted to focus student interactions on different topics) or through orchestrational support (e.g., where physical elements of the space, like projected displays, help to guide or coordinate student movements, collaborations, or activities). S3 was also developed to add a level of intelligence to classrooms or other learning environment, including real-time data mining and computation performed by intelligent agents to support the orchestration of inquiry scripts. In addition, we are also investigating the role of ambient displays

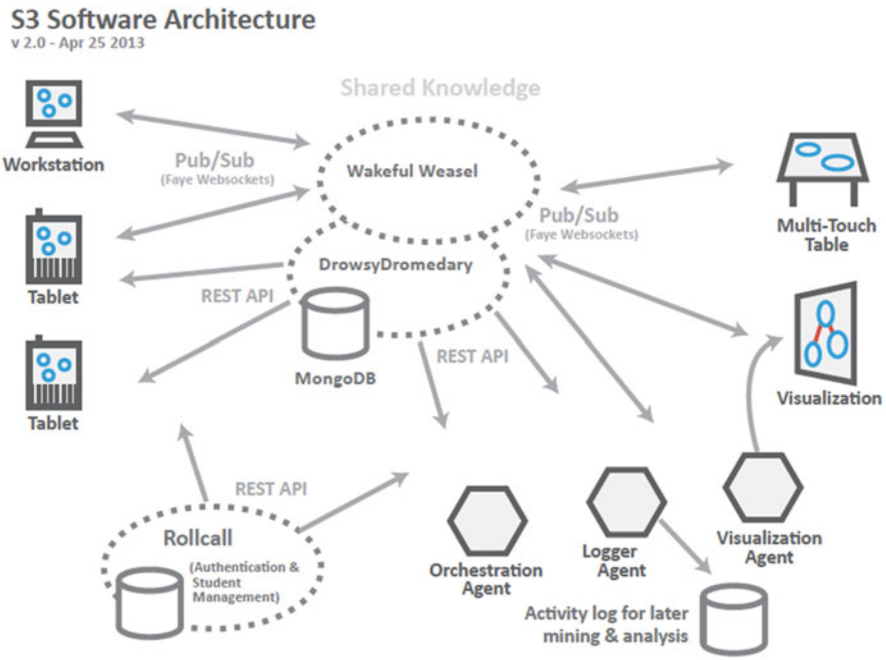

Fig. 12.2 SAIL smart space (S3) systems architecture, showing the use of direct WebSocket messaging to enable communications among any element of the environment, a persistent, non-relational (no SQL) database (MongoDB) and intelligent agents

of information, within the physical environment, as a means of providing "peripheral" guidance or feedback (Alavi et al. 2009).

At present, S3 includes a set of core technologies: (1) a portal for student accounts and software application management, (2) an intelligent agent framework for data mining and tracking of interactions in real time, (3) a central database that houses the designed curriculum and the products of student interactions, and (4) a visualization layer that controls how materials are presented to students (see Fig. 12.2). Our goal in developing S3 was to support a broad program of research on collaborative inquiry, allowing for more rapid development of learning materials and environments. While it is not designed as an off-the-shelf solution, S3 is offered as an open-source framework, in hopes of promoting wider access to such functionality and growing a community of developers within the learning sciences (Slotta et al. 2012).

This paper presents an application of KCI and S3 to enable seamless learning across contexts, including user-contributed, tagged, and coedited materials, and a role for intelligent agents in coordinating a complex sequence of student activities. Working closely with a high school physics teacher, we designed and developed a 12-week physics curriculum where students engaged in learning activities across several contexts: (1) their classroom, (2) their homes, (3) field observations, and

(4) a “smartroom” (different from their classroom) where they engaged in carefully scripted interactions with an array of media and materials. We investigated an orchestrational framework that included the use of social tagging and metadata, as well as intelligent agents to support the enactment of collaborative inquiry scripts. The following sections will detail the design of the curriculum, the relevant technologies, and our enactment with two sections of a high school physics course.

Methods

Throughout the design process we employed a codesign methodology (Roschelle et al. 2006) working closely with a high school physics teacher to ensure that he was an active voice in the technology design and that the designed intervention “fits” his goals for the students and expectations for student learning. Because the research was situated within a real class (rather than a canned lab setting), a design-based research approach was implemented in order to respond to the multitude of variables present during its enactment (Wang and Hannafin 2005). Generally, design-based research does not attempt to validate a particular curriculum; rather it strives to advance a set of theories on learning that transcend the particulars in which they were enacted (Barab and Squire 2004). As such a major outcome of this research was the design of the curriculum and supporting technologies themselves. In order to evaluate the enacted design, we used a mixed methods approach in order to triangulate the data and get a more complete picture (Johnson et al. 2007; Mason 2006). Sources included pre- and post-interviews with students and the teacher, server logs, the user-contributed artifacts, and video and audio recordings during the culminating activity.

Physics Learning Across Contexts and Environments (PLACE)

In order to investigate the role that such technologies could play in supporting cross-context learning, we needed to develop a carefully designed curriculum that leveraged student-contributed content and included a meaningful role for intelligent agents and data mining. We work closely with the high school physics teacher to develop a curriculum that implemented KCI, including collaborative and collective forms of inquiry and adding a level of critical reflection to the teacher’s previous approach. Two main goals were identified by the teacher: First, he wanted to help students to recognize “physics in their everyday lives” and then bring this view of physics back into the traditional classroom setting. Second, he wanted to design some way for students to develop a coherent understanding of the underlying principles of the course, including the connections among those physics principles (i.e.,

Table 12.1 Grade 11 fundamental principles for kinematics, force and motion, and work, energy, and power

Vectors	Acceleration	Fnet=0	Kinetic energy
Newton's First Law	Uniform motion	Fnet=constant (nonzero)	Potential energy
Newton's Second Law	Kinetic friction	Fnet=nonconstant	Conservation of energy
Newton's Third Law	Static friction		

to "see that all the principles are tied together"). We began by generating a list of fourteen principles (Table 12.1) that covered the first three units of the course: (1) kinematics, (2) forces and motion, and (3) work, energy, and power. Following the work of Chi, Feltovich, and Glasser (1981), we wondered if, by engaging students in principle-based classification of physics phenomena and problems, we could help them achieve a greater level of expertise.

In order to achieve these goals, we developed a script that engaged students in capturing examples of physics in the world around them (either through videos, pictures, or text), which they uploaded to the classroom database, "tagged" with any of the principles they felt to be applicable, with a written explanation for their choice of tags. The wider community of students was encouraged to respond to these user-contributed artifacts: debating tags or explanations, voting, and adding new tags – with the stated aim of developing consensus about each item. To support this process, we developed a micro-script that required students to complete three steps of (1) voting on existing tags and/or adding a new tag, (2) voting on the contributions of their peers, and (3) adding a reflection or rationale of their own. This was designed to ensure that students covered three key aspects (focus on the principles, reflecting on the work of their peers, and adding their own thinking). As part of the script, in order to ensure that all the principles were covered and to encourage students to become experts in particular principles, we assigned each student to an "expert group" in which they were assigned a subset of the principles (e.g., Newton's First Law, vectors, and potential energy) for which they were responsible to keep updated (i.e., to make sure all relevant items had been tagged and add a comment where they felt the principles had been wrongly tagged).

For each of the three units within the curriculum, students were tasked with uploading at least one example and to commenting on at least two of their peers' submissions (with a focus on their expert principles). At the end of each unit, the teacher selected some examples to discuss with the class and had the students look over examples that had been tagged with their principles, to add to their discussion. Students also uploaded results from their in-class laboratory experiments to the knowledge base, tagging their reports with principles and adding reflections on their methodologies – which other students were also free to critique. For homework, the teacher provided multiple-choice problems, which students solved using a script similar to the one used for their contributions: tag, answer, and reflect on the problem. All student contributions went into a collective knowledge base, which itself served as a basis for various further activities. For example, students were asked to develop "challenge homework problems" for their peers, using examples drawn from the knowledge base. Intelligent agents mined the knowledge base to retrieve principle-tagged problems during the smart classroom activity.

The teacher's role in the curriculum was also scripted, in the sense that he was expected to upload regular homework problems, review and assess student answers, and adjust class lessons accordingly. He was also expected to review student contributions, to find examples or interesting discourse for use during in-class discussions. Finally, his role was tightly scripted in the smartroom activity, where he had consequential roles of approving students when they had gotten to a certain point, providing feedback if he did not approve and leading whole class discussions.

S3 Supports for PLACE: Learning Across Contexts

To support student interactions at home, in their neighborhoods, and in the classroom, we needed a technology infrastructure that supported student activity, from completing homework problems, to uploading examples, to tagging and discussing peers' contributions. S3 supported our development of two complementary systems: *PLACE.Web* (Physics Learning Across Contexts and Environments), a collaborative social network, focused on the domain of physics, where students contribute content, engage with the work of their peers, and complete tasks assigned by the teacher, and *PLACE.neo*, a smartroom environment that orchestrated the activity, making use of the PLACE.Web content. Both PLACE.web and PLACE.neo employed elements of S3, including Rollcall, a user portal that provided each student a personal profile and nickname and lets them personalize their identity within the community. The use of Rollcall also allowed S3 intelligent agents to personalize the kinds of information visible to each student, the materials they were actively provided, and their group assignments in the culminating activity.

PLACE.web

The PLACE.web learning environment supported five different interaction spaces for the students: (1) the student status page, (2) the contribution upload page, (3) the user contribution discussion pages, (4) the assigned homework pages, (5) and the "associative web" – a semantically aggregated visualization of the entire community knowledge base.

- *The student status page* – This was the first page that students saw when logging into PLACE.web and was broken into several distinct information spaces to give the student a quick overview of their contributions and the state of the overall class activity (Fig. 12.3). The goal of this page orients students' personal place within the knowledge community and provides insight into possible avenues for

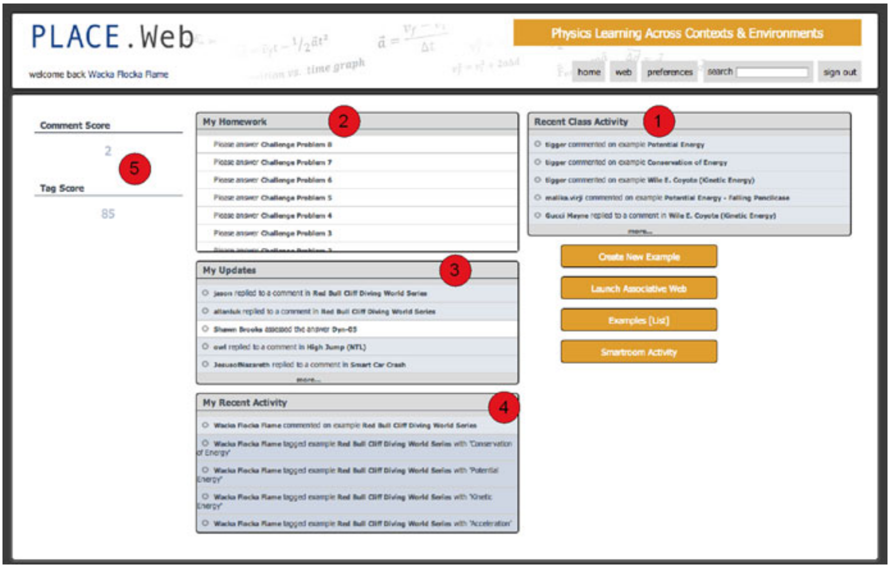

welcome back Wacka Flocka Flame

Comment Score: 2

Tag Score: 85

My Homework

Please answer Challenge Problem 8
Please answer Challenge Problem 7
Please answer Challenge Problem 6
Please answer Challenge Problem 5
Please answer Challenge Problem 4
Please answer Challenge Problem 3

My Updates

○ jason replied to a comment in Red Bull Cliff Diving World Series
○ allanluk replied to a comment in Red Bull Cliff Diving World Series
○ Shawn Brooks associated the answer Dyn-03
○ owl replied to a comment in High Jump (NTL)
○ JesusofNazareth replied to a comment in Smart Car Crash

My Recent Activity

○ Wacka Flocka Flame commented on example Red Bull Cliff Diving World Series
○ Wacka Flocka Flame tagged example Red Bull Cliff Diving World Series with 'Conservation of Energy'
○ Wacka Flocka Flame tagged example Red Bull Cliff Diving World Series with 'Potential Energy'
○ Wacka Flocka Flame tagged example Red Bull Cliff Diving World Series with 'Kinetic Energy'
○ Wacka Flocka Flame tagged example Red Bull Cliff Diving World Series with 'Acceleration'

Recent Class Activity

○ tigger commented on example Potential Energy
○ tigger commented on example Conservation of Energy
○ tigger commented on example Wile E. Coyote (Kinetic Energy)
○ malika.virji commented on example Potential Energy - Falling Pencilcase
○ Gucci Mayne replied to a comment in Wile E. Coyote (Kinetic Energy)

Fig. 12.3 The student status page had several informational streams to help students orient themselves within the knowledge community and manage their orchestrational load through (*1*) the recent class activity, (*2*) my homework, (*3*) my updates, (*4*) my recent activity, and (*5*) comment and tag scores

action when the teacher was not around to directly provide instruction or guidance. As such this page was one of the focal points of the informal student learning activities.

- The status page showed a news feed of the whole class' contributions ("Recent Class Activity"), giving the student a sense of the overall class activity and a means to jump to any particular artifact or comment they might have found interesting. The other feeds were personalized to the individual student: The "My Homework" feed showed students any tasks assigned to them by the teacher, which would automatically disappear once a student had completed the task; the "My Updates" feed showed the student any actions that other members of the community made on any of his or her contributions (e.g., commented on one of their examples, agreed or disagreed with one of their tags), providing students an active connection to the knowledge community and their place within it; the "My Recent Activity" feed tracked all the actions by the individual student, giving a means of tracking his or her own contributions to the community and quickly jumping to a space of interest (i.e., where he or she is involved in discourse). On the left side of the status page, each student saw a "Comment Score" and a "Tag Score," which tracked the total votes students had received from their peers for their contributions. This provided a means of motivating students to produce "high-level" contributions.

- *The contribution upload page* – This is where students uploaded their contributions (video, picture, or narrative) to the shared knowledge base. In addition to their uploaded media, as part of the scripted interactions, students were required to also assign tags and a rationale of their physics thinking to the contribution. The contribution upload page was designed to be as device agnostic as possible to allow students to upload and create content in a broad range of contexts (at home, in their neighborhoods, at school). We aimed to facilitate a level of mobile integration to PLACE.web, and students using Android devices could upload media directly from their device to PLACE.web, allowing them to capture physics "on the go." Students using iOS (iPhone, iPad) needed to first transfer their media to another computer before contributing it to the knowledge base.
- *The discussion pages* – The discussion pages (Fig. 12.4) in PLACE.web were designed to allow students to engage in discussion and debate and vote on the principles tagged to the contribution. These interactions took the form of threaded discussions, including aggregated votes for each of the principles. The contribution pages were used widely throughout the script, as students would engage in these spaces both at home and during scripted in-class sessions. These pages were also designed to be as device agnostic as possible so that students could access and contribute to them from any major browser, as well as from both Android and Apple devices.
- *The assigned homework pages* – These pages were teacher created and were centered on multiple-choice homework problems. The scripted interface was similar to that of contributions where students had to tag and provide a rationale

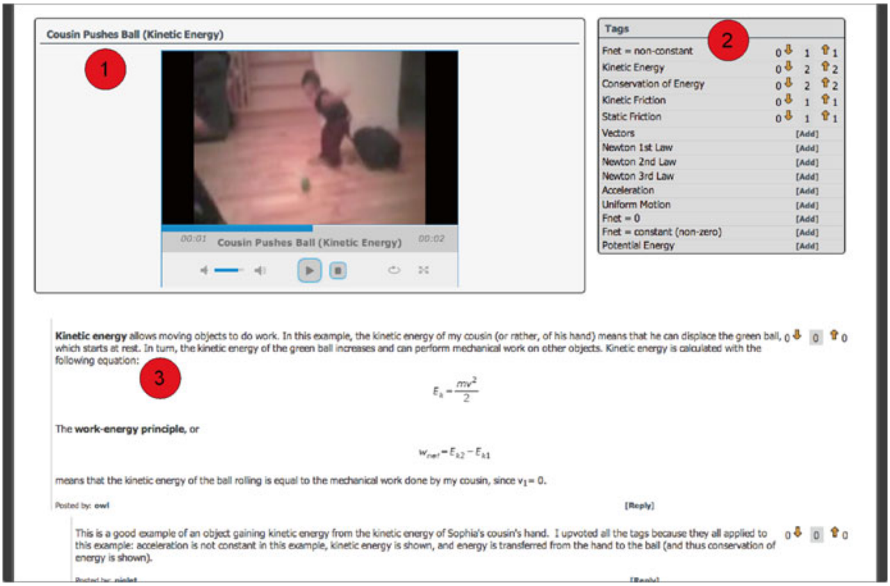

Tags			
Fnet = non-constant	0 ⇩	1	⇧ 1
Kinetic Energy	0 ⇩	2	⇧ 2
Conservation of Energy	0 ⇩	2	⇧ 2
Kinetic Friction	0 ⇩	1	⇧ 1
Static Friction	0 ⇩	1	⇧ 1
Vectors		[Add]	
Newton 1st Law		[Add]	
Newton 2nd Law		[Add]	
Newton 3rd Law		[Add]	
Acceleration		[Add]	
Uniform Motion		[Add]	
Fnet = 0		[Add]	
Fnet = constant (non-zero)		[Add]	
Potential Energy		[Add]	

Kinetic energy allows moving objects to do work. In this example, the kinetic energy of my cousin (or rather, of his hand) means that he can displace the green ball, which starts at rest. In turn, the kinetic energy of the green ball increases and can perform mechanical work on other objects. Kinetic energy is calculated with the following equation: 0 ⇩ 0 ⇧ 0

$$E_k = \frac{mv^2}{2}$$

The **work-energy principle**, or

$$W_{net} = E_{k2} - E_{k1}$$

means that the kinetic energy of the ball rolling is equal to the mechanical work done by my cousin, since $v_1 = 0$.

Posted by: owl [Reply]

This is a good example of an object gaining kinetic energy from the kinetic energy of Sophia's cousin's hand. I upvoted all the tags because they all applied to this example: acceleration is not constant in this example, kinetic energy is shown, and energy is transferred from the hand to the ball (and thus conservation of energy is shown). 0 ⇩ 0 ⇧ 0

Fig. 12.4 An example of a contribution discussion page with (*1*) a student-uploaded video, (*2*) student-submitted principle tags and voting, and (*3*) threaded student discourse

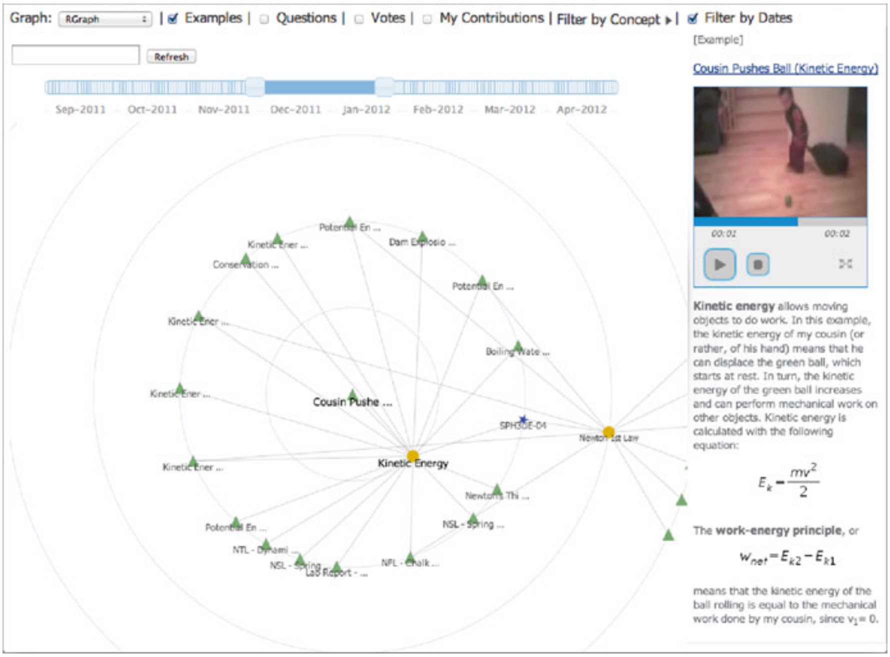

[Example]

Cousin Pushes Ball (Kinetic Energy)

Kinetic energy allows moving objects to do work. In this example, the kinetic energy of my cousin (or rather, of his hand) means that he can displace the green ball, which starts at rest. In turn, the kinetic energy of the green ball increases and can perform mechanical work on other objects. Kinetic energy is calculated with the following equation:

$$E_k = \frac{mv^2}{2}$$

The **work-energy principle**, or

$$W_{net} = E_{k2} - E_{k1}$$

means that the kinetic energy of the ball rolling is equal to the mechanical work done by my cousin, since $v_1 = 0$.

Fig. 12.5 Associative web, showing filtered view of the principles "kinetic energy" and "Newton's First Law," with examples from student contributions

in addition to their answer; however, with the homework problems the contributions of their peers were not shown to students. As with the discussion pages, students could access the pages from any major browser or mobile device.

- *The Associative Web* – The Associative Web (Fig. 12.5) was an interactive, filterable visualization that used the principle tag metadata to semantically connect all the contributions of the knowledge community. The Associative Web was used primarily during in-class activities in which students were tasked with finding examples that shared principles with their assigned expert group and when students were finding examples to scaffold the creation of their challenge problems. The teacher also used the Associative Web as a tool for in-class discussion by examining the clustering of student contributions as a way of finding similarities between seemingly disparate physics examples.

The teacher was also provided with tools to manage his orchestrational load, including a front status page that showed a similar set of feeds to those seen by the students. Akin to the student contribution upload page, the teacher was provided an authoring page that allowed him to create multiple-choice problems in a few short clicks. The teacher was also provided with two additional tools to give him insight into the class for adjusting the script based on understanding the class' emergent knowledge.

- *Built-in assessment* – The teacher was provided with a customized assessment tool on each contribution or homework page, allowing him to provide students with a mark (from 1 to 4) and personalized feedback. The assessment tool also allowed the teacher to write himself personal notes based on the student assessment to review toward adjusting upcoming lessons.
- *Individual student reports* – The teacher was provided a single page that provided detailed information on each student's activity on the site including links to his or her individual contributions and their total and average marks from his assessments of their work.

PLACE.neo: Leveraging Student-Contributed Materials and Tagging for New Learning Contexts

The goal of the culminating activity, following the KCI model, was for students to make use of their co-constructed knowledge base in the context of some final inquiry activity. Another important goal of this research was to investigate the technology infrastructure of S3, including some strong role for intelligent agents and real-time data mining. After many design discussions, we arrived at a challenging task that involved analyzing the physics of Hollywood movie clips, including setting up physics problems to test their validity. This culminating activity involved three micro-scripts that spanned home, a traditional class setting and a smart classroom, and relied heavily on S3 agents to coordinate the distribution of materials, roles, and tasks. At home, students were tasked with looking at a collection of the problems they had been assigned during the preceding 12-weeks (including their contributed challenge problems and new problems developed by the teacher), verifying their tagging of relevant physics principles, and adding equations that might be used to solve the problems. In class, students worked in small groups, using tablet computers to reach consensus on a refined "final set" of the tags and equations for each problem. These tagged problems, principles, and equations were thus processed by students from their collective knowledge base, to be used as a prepared set of materials within the final smart classroom script, where intelligent agents would access and distribute them.

Once entering the smart classroom, students were engaged in solving a series of ill-structured physics problems using Hollywood movie clips as the domain for their investigations (e.g., could Iron Man survive a fall to earth, as depicted in the movie?). Four videos were presented to the students, each at a distinct physical location within the room (Fig. 12.6). The students were engaged collectively, working as a whole group of 12–16, as well as collaboratively, in various small group configurations as commanded by the S3 intelligent agents. Agents made grouping decisions according to predefined scripting criteria, relating to the students' use of principles within an initial tagging activity and to the need to regroup students with peers they had not worked with yet. The smartroom script was broken up into four

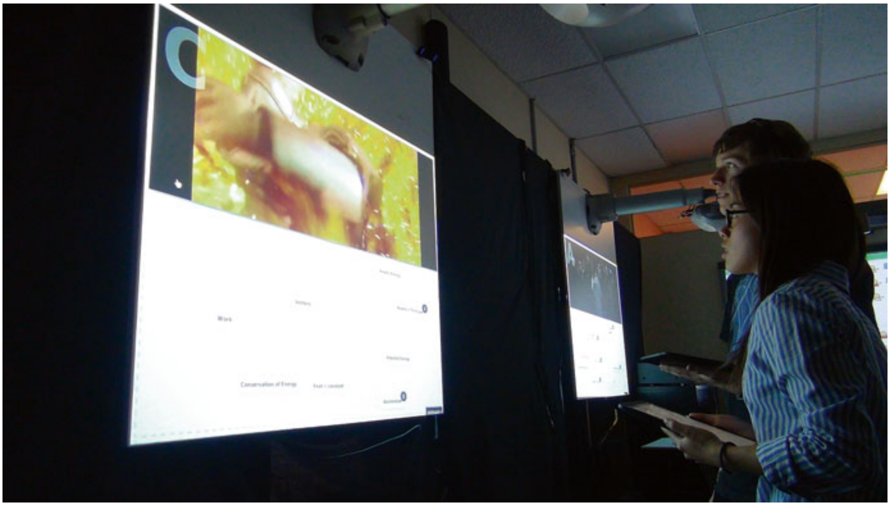

Fig. 12.6 Students engaging with the interactive displays and individual tablets in the smart classroom

different steps: (1) Principle Tagging, (2) Principle Negotiation and Problem Assignment, (3) Equation Assignment and Assumption and Variable Development, and (4) Solving and Recording. In each step students moved from one video to another, completing a set of collective and collaborative tasks that built upon the emerging knowledge base, using tablets and large-format interactive displays.

Technology Implementation of PLACE.neo

The script for the culminating activity relied heavily on the S3 agent framework in order to coordinate the complex distribution of materials, roles, and tasks. To scaffold the different contexts (at home, in class, in the smart classroom) and interactions (individual, cooperative, collaborative), we developed specific technology supports for each stage of the activity, in order to connect student activities with the knowledge base and achieve the overarching pedagogical goals of the script (i.e., solving the Hollywood video problems).

In order to facilitate the at-home portion of the script, and capitalize on the students' familiarity with the platform, we implemented this first stage using PLACE.web, adding a new icon to the existing student status page for students to access the activity. Drawing on the metadata that indicated each student's assigned content expertise, PLACE assigned each student a specific subset of the problems to tag with principles and equations. The use of the metadata allowed us to customize the problem sets seen by each student to ensure that every problem was "covered" by all fourteen principles.

During the in-class portion of the culminating activity, we developed a context-specific tablet application that connected students to their peers in real time, using the aggregated products of the previous at-home stage. Once again, we used the students' expertise metadata to group them, sending each student's group assignments to his or her tablet. The goal of this activity was for students to achieve consensus about the principles and equations that had been assigned to each problem in the corpus. The group would be assigned one of the problems, with each student seeing the problem and its various tags on his or her tablet, and asked to agree or disagree. In order to ensure consensus was achieved, we employed a *Consensus Agent,* which required all students within a group to have the same choices on their tablets before moving to the next task. Students could see the work of their group members in real time, reflected on their own tablets, which facilitated face-to-face discussions. We distributed all the problems to the groups using an S3 *Bucket Agent,* where a group was provided with a new problem once it had finished its existing one, and when all problems were gone (i.e., the "bucket" of problems was empty), they received a "please wait" message. In this way, groups who worked faster or who had received easier problems were given more problems, such that a large number of problems were addressed in an efficient, distributed fashion.

For the smart classroom stage of the script, we developed a script, which took advantage of the physical and collaborative affordances of the smart classroom, including large projected displays accompanying each video station and individual tablet computers to support students as they performed activities. The students' tablets coordinated all activities, populated by intelligent agents that drew content from the products of the in-class activity. Students worked in small groups, with the products of their individual tablet interactions aggregated and broadcast to the large group display, which then led to further collaborative knowledge building tasks. S3 agents queried metadata to provide students with context-specific tasks and materials, drawn from the corpus of student-contributed and student-tagged materials from earlier activities. The smart classroom script consisted of 4 steps (see Fig. 12.8), described below.

In step one, each student received a set of three or four principles (i.e., out of the 14) on their tablet, determined by querying that student's prior expertise groups. The students were asked to go to one video at a time and to "swipe" any of their four principles that they found relevant to the video onto the large display at that station. After four 2-min intervals, all students had tagged each of the four videos with any of the principles that were relevant. Because each principle had been assigned to at least two students, there were multiple instances of the principles on the boards (see Fig. 12.9).

In step two, students were assigned, by an S3 *Student Sorting Agent*, to one of the video boards, according to where they had swiped the most principle tags (while still evenly distributing the students around the room). The student tablet provided the ID of the video to which he or she was assigned (e.g., "A," "B," "C," or "D") and walked over to that video station. Once all students had arrived at their assigned stations, the teacher "advanced" the script using his tablet, and students received their task: They first negotiated, with the aid of a *Consensus Agent*, the

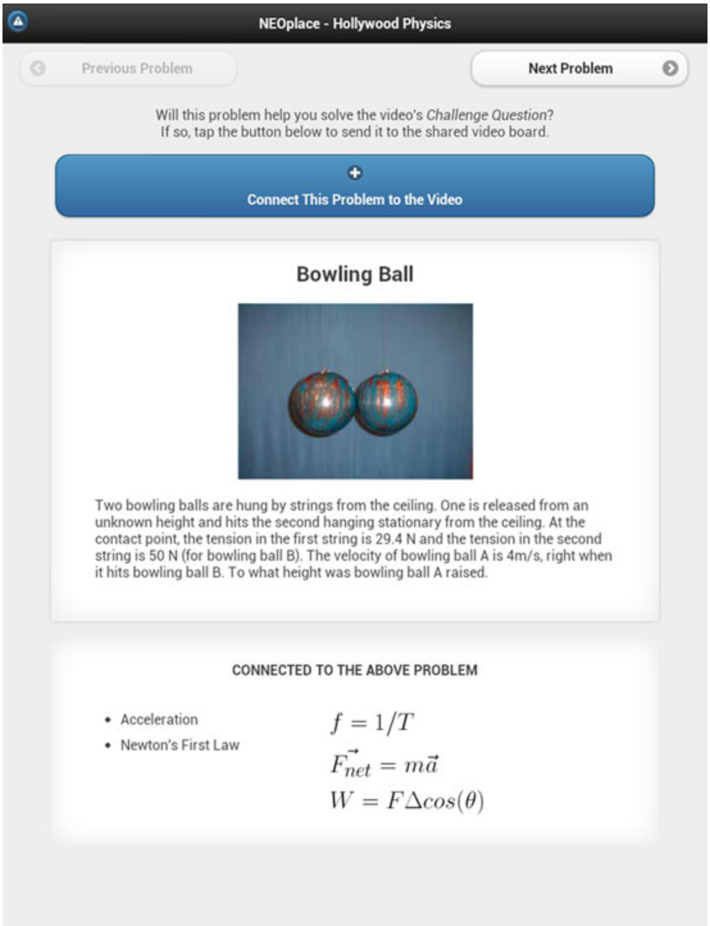

Will this problem help you solve the video's *Challenge Question*?
If so, tap the button below to send it to the shared video board.

Bowling Ball

Two bowling balls are hung by strings from the ceiling. One is released from an unknown height and hits the second hanging stationary from the ceiling. At the contact point, the tension in the first string is 29.4 N and the tension in the second string is 50 N (for bowling ball B). The velocity of bowling ball A is 4m/s, right when it hits bowling ball B. To what height was bowling ball A raised.

CONNECTED TO THE ABOVE PROBLEM

- Acceleration
- Newton's First Law

$$f = 1/T$$

$$\vec{F}_{net} = m\vec{a}$$

$$W = F\Delta cos(\theta)$$

Fig. 12.7 Individual student tablet screen from the smart classroom activity's problem assignment task (step 2)

final principles for their video. Then, another S3 Bucket Agent retrieved the physics problems that had been tagged with those principles in the previous (in-class) activities and distributed them to the individual students within the group, who made simple "yes or no" decisions about whether the problem might be an interesting model for how to set up the video, for solving in problem form (Fig. 12.7 above). Unlike the *Bucket Agent* in the in-class activity (whose goal was to get students through the task as efficiently as possible), this agent aimed to get all the students in a group involved in idea promotion and negotiation. As such, each member in the group received an equal but unique set of items that were semantically connected to their video by the S3 agent. As part of the script each student had to promote at least one problem to the negotiation board from their set (Fig. 12.10), encouraging each student to take an active role in setting up the

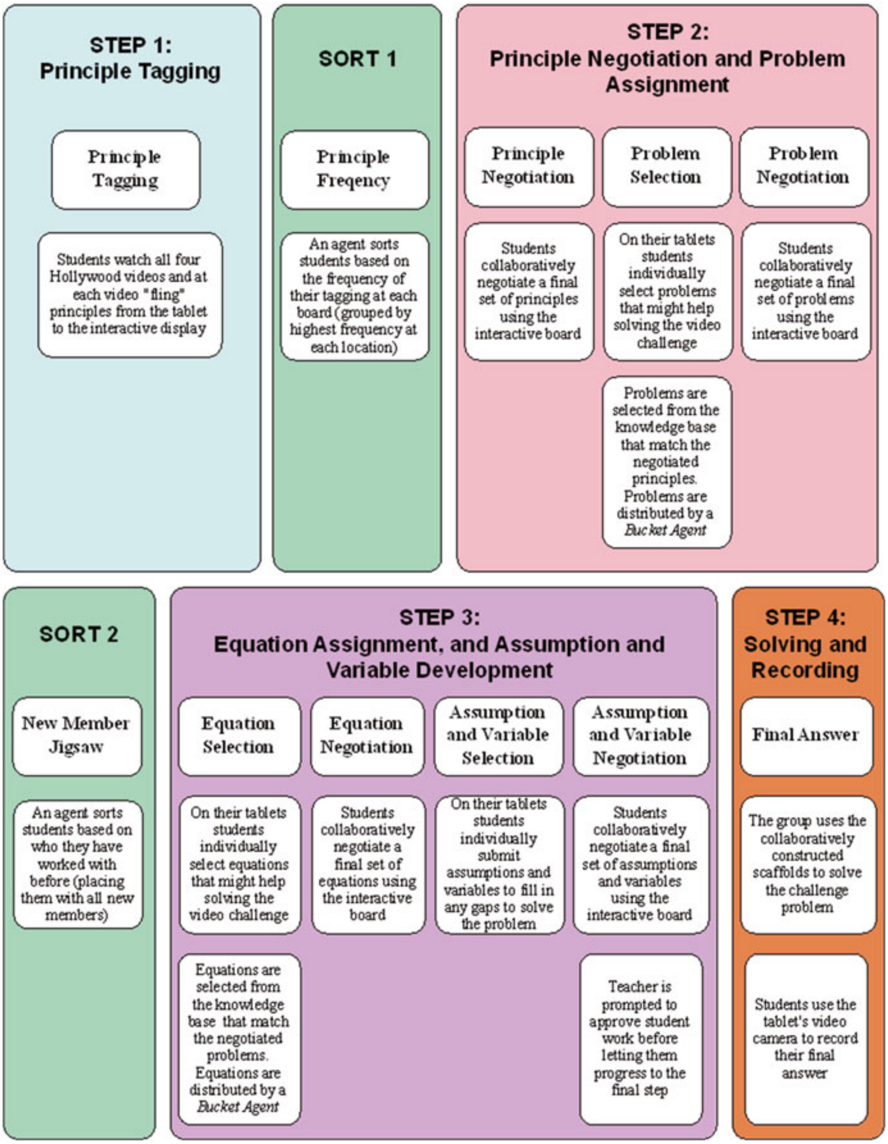

Fig. 12.8 The smart classroom Hollywood physics script involved four distinct steps

problem. Additionally the movement between the "private" space of the tablet and the public and collaborative space on the interactive walls aimed to have students work in multimodal contexts within the activity.

In step three, a *Student Sorting Agent* reassigned students to new video stations, based on a criterion for grouping students who had not worked together in any

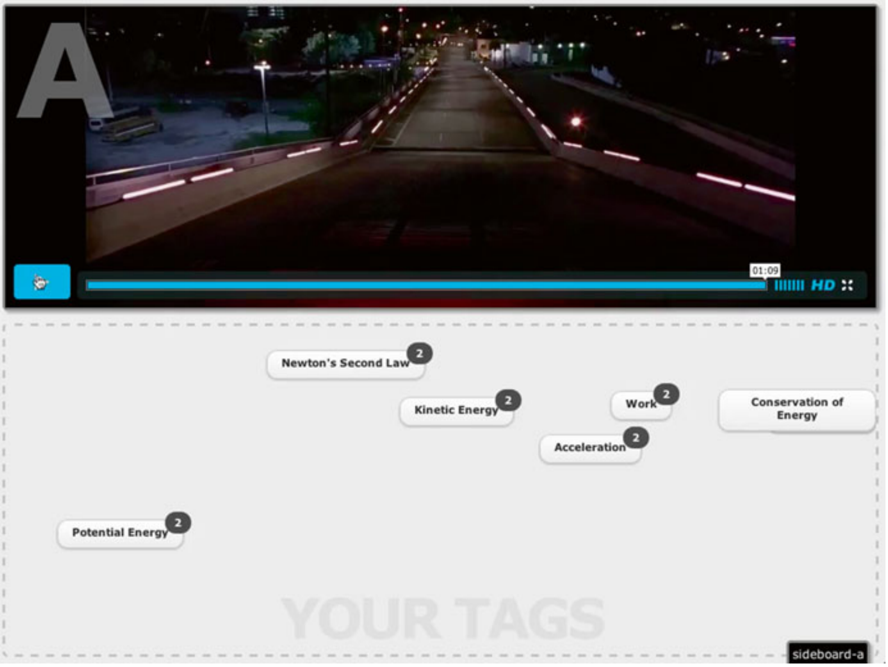

Fig. 12.9 Interactive board during step 1 – the principle tagging phase. The numbers indicate how many times the video was tagged with that principle (e.g., two students tagged the video with "Work")

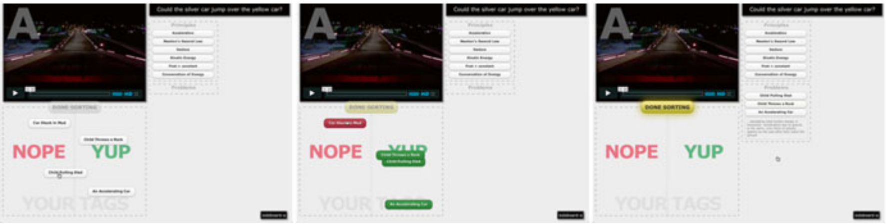

Fig. 12.10 Shows the three phases of the Problem Selection task (step 2), where students (*1*) submit problems from their tablets to the interactive board, (*2*) negotiate which problems to keep or discard by dragging them to the "Yep" or "Nope" zone of the negotiation space, and (*3*) after negotiation the final set appears on the *right*

previous step. A *Bucket Agent*, similar to the one employed in step two, distributed the problems to students at each board and showed them the equations connected to the problem during the in-class portion of the script. Students promoted those equations they felt might help in solving the challenge question to the shared display and negotiated a “final set” which was again facilitated by the Consensus Agent. Group members then individually came up with assumptions and variables to fill in any information “gaps” and engaged in the negotiation and consensus script to produce a final set. Unlike with the other negotiation and consensus tasks, when a group submitted a final set of assumptions and variables, the teacher was alerted on his orchestration tablet to review students’ work and either approve it or to send them back to refine their submission.

In step four, the final step, student groups used the collaboratively constructed scaffolds on the interactive whiteboards for support and with pen and paper solved their challenge problem and recorded their final answer as a video narrative using the tablet’s built-in video camera.

A critical part of this design was that in order to make these complex orchestrations occur and draw materials that had been tagged during the previous in-class stage, the S3 agents needed to be able to respond to the emergent conditions in the class. The S3 agents could not know what tags the students would choose in step one (which would determine the problems selected for step 2 or the problems selected from step 2 for the equations for step 3), and therefore had to be developed as adaptive scaffolds responding to the real-time products of the class’ knowledge construction.

Enactment of the PLACE Script

PLACE was implemented with 2 sections of grade 11 physics ($n=22$, $n=22$) in an urban high school. Over the 12-week curriculum, the students were actively involved in the development of artifacts and in the discussion around the physics principles connected to them. Students regularly uploaded examples to the database and engaged in discussion around their physics principles. Below we discuss students’ contributions across different contexts and their subsequent reuse in class and during the smart classroom activity.

Student-Contributed Content

In total 169 student examples were created, and 635 total student discussion notes were contributed around those examples. Students also attached 1,066 principle tags to the contributed examples and cast 2,641 votes on those tags. Although the designed script required students to upload at least one example in each of the three units (3 contributions in total), students on average submitted 3.84 examples to the knowledge base (excluding the challenge problems), which seems to point to active community engagement.

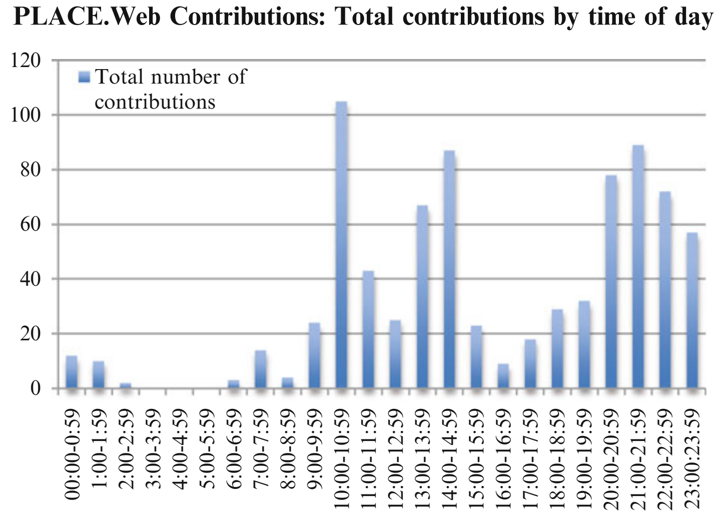

Fig. 12.11 Graph shows the total number of student contributions to PLACE.web distributed by time of day

During the enactment of the script, students were actively engaged in school, at home, and in their neighborhoods. An examination of the time of day at which students contributed to the knowledge base shows that uploads or comments were made within PLACE at nearly every point of the day (the only exception being between 3 am and 6 am, see Fig. 12.11). This highlights the ability of PLACE.web to seamlessly connect students within their overall community whenever they felt the desire to take part, with 46.58 % of contributions taking place during school hours (9 am–4 pm) and 53.42 % of the contributions taking place outside of school hours (4 pm–9 am). Interestingly, nearly 2 % of the overall interactions took place during students' lunchtime (12:30 pm–1:15 am), which indicates both the interest and ability to access PLACE outside of traditional in-class hours, even while still in the formal school setting. The teacher involved in the study noted that several times students came up to him in the hall with their mobile devices, to bring up a homework question or a peer's example, and asked his thoughts about their response. He stated that he was amazed not only at their interest but also their ability to have the content "at their fingertips."

An examination of the types of student-contributed content also highlights the seamless nature of PLACE. Student-contributed examples included videos of friends at a track meet, a subway arriving at a station, a student pushing a friend into a pool, a student's young cousin rolling a ball in their house, and a pair of students rolling two different-sized objects down the school's hallway. All of these examples point not only to the ability of PLACE to capture moments of student insight but also, perhaps more critically, that the curriculum, and PLACE as a support for the curriculum, got the students actively seeking out, capturing, and questioning physics in their everyday lives.

Using Peers' Contributions: The Challenge Problem Script

Working collaboratively in groups of three to four in the classroom, students were tasked to create "challenge problems" that would be solved by their peers, drawing from the wider knowledge base of peer-contributed examples. This script was seen to engage students and leverage their collective knowledge base, leading to the development of further materials for peer engagement and investigation. In total, 13 challenge problems were developed by students, each of which referred to, on average, 2.23 examples from the knowledge base. The Associative Web was employed as an in-class tool to help students find examples that matched their expertise groups and supported their creation of a challenge problem. In post-activity questionnaires, students indicated that they found the Associative Web very useful for filtering the overall knowledge base and to find artifacts that matched their individual search criteria, noting that "*the examples about each concept were easily identified and similar examples were grouped together,*" and "*the associative web made it clear what examples are related to our concepts, because you could see what example was related to more than one of the concepts, and it's easy to browse through multiple areas.*"

Culminating Activity: Scripting and Orchestration Across Contexts

During the culminating smart classroom activity, students were able to access, contribute to, and use the knowledge base at home, in class, and in the smart classroom. This activity was an important test of the capabilities of S3 to support seamless orchestration of learning activities, and the use of intelligent agents was central to our success. The next three sections address how S3 agents supported learning in each of the three contexts.

Agent Orchestration of the At-Home Activity

In the at-home portion of the culminating activity, students were scaffolded in answering a subset of the homework problems, depending on what "expertise groups" they had been assigned to in previous units. S3 agents were employed to ensure that each problem ($n = 30$) was received by students who represented all fourteen principles. This was successfully achieved, ensuring that every problem in the corpus was reviewed by the knowledge community in terms of every principle.

Agent Orchestration of the In-Class Activity

During the in-class portion of the activity, S3 agents successfully grouped students and facilitated their consensus building on all of the homework problems. Of particular interest within the in-class activity was the effectiveness of the *Bucket Agent*

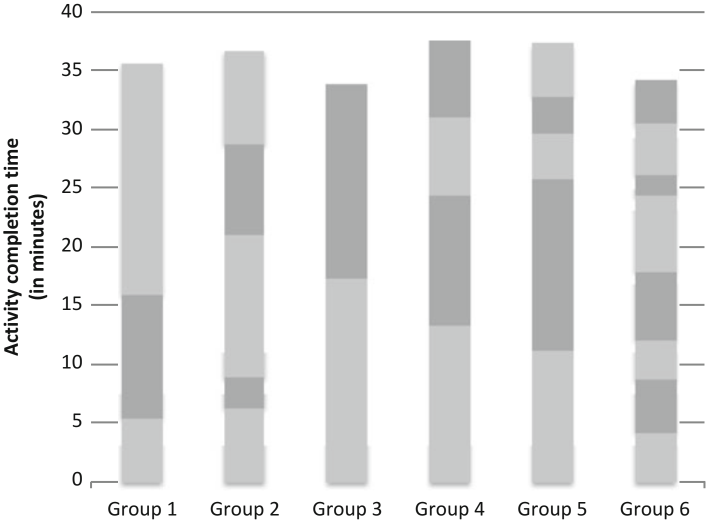

Fig. 12.12 *Shaded bars* show the number of problems sent to each group by the *Bucket Agent* and how long the students spent on each problem. For example, group 3 took a long time on both of its problems, so they only received 2

in orchestrating the real-time distribution of the problems to the individual groups. Given the time constraints in the classroom – only 60 min, which included all the kids arriving in class, taking their seats, and the researchers distributing the tablets and explaining the activity *before* starting – it was imperative that the problems were distributed as efficiently as possible. The *Bucket Agent* regulated the distribution of problems in such a way that every group completed their assigned problems in less than 40 min and within 3 min of each other (Fig. 12.12), minimizing the variance between time on task for groups.

Agent Orchestration of the Smart Classroom Activity

Within the smart classroom portion of the culminating activity, the S3 agents successfully responded to emergent properties of student interactions to supply them with semantically relevant artifacts, drawn from the in-class activity. During step 2 of the smartroom activity (the "Problem Assignment" step), students were given problems, drawn by agents from the knowledge base, whose principles matched those that had been assigned to their video clip. The S3 agents connected, on average, 23 problems to each video, of which students agreed (voted "yes") to an average of 3.4 problems, which were negotiated down (during the whiteboard consensus phase) to an average of 2.6 problems. During step 3 ("Equation Assignment"), S3 agents were able to draw, from the knowledge base, the equations that had been assigned to those problems, to serve as resources for students in setting up their solutions to the video clip challenges. From these agent-filtered equations, students

Table 12.2 Student tagging frequencies and sorting agent assigned boards for step two and step three

Students	# of tags by student Board A	# of tags by student Board B	# of tags by student Board C	# of tags by student Board D	First sort: sent to board	Second sort: sent to board	Sorted to new board?	Sorted with new team members?
Alice	4	3	3	4	A	B	Y	Y
Pearl	3	0	3	2	A	C	Y	Y
Jason	4	3	4	4	B	C	Y	Y
Rob	0	3	3	1	B	D	Y	Y
Desi	3	2	3	0	C	D	Y	Y
Raffi	0	2	2	2	C	A	Y	Y
Becky	2	2	3	3	D	A	Y	Y
Sun	2	2	0	2	D	B	Y	Y

recommended an average of 4.9 equations, which were negotiated down to an average of 4.3 equations, during the whiteboard negotiation phase.

In between steps of the smart classroom activity, the *Student Sorting* Agent was able to successfully sort students into groups based on the number of principles they had signed to each video (step 2) and ensuring they were working with new groupmates (step 3). We approached this challenge by having the agent build a table of student interactions (similar to Table 12.2, above), which was used in a cascading fashion to assign one student to board A based on their frequency of principles, then one to boards B, C, and D in order, before repeating this process until all students were sorted. Jason was assigned to board B and not A, C, or D because the agent had already placed Alice at board A, and Jason had the most tags when the agent went looking for a board B student (i.e., for the second assignment by the agent's algorithm).

Solving the Hollywood Physics Problems

In the final step of the smart classroom activity, it is important to note that students were successful in setting up and solving the Hollywood film clips, using the assumptions and equations that had been generated from previous steps. Every group succeeded, in the time allowed, in generating a written solution to the problem and creating a short video where they explained their solution.

By looking at the final state of the collaboratively built knowledge on the interactive whiteboards and comparing it to the elements (such as the assumptions and variables and equations) used by the students in solving the problem (Fig. 12.13), we can begin to see how the interactive board was useful for scaffolding the students' problem solving. The exit interviews with students supported the visual evidence of the value of the boards, the user-contributed content in this scaffolding

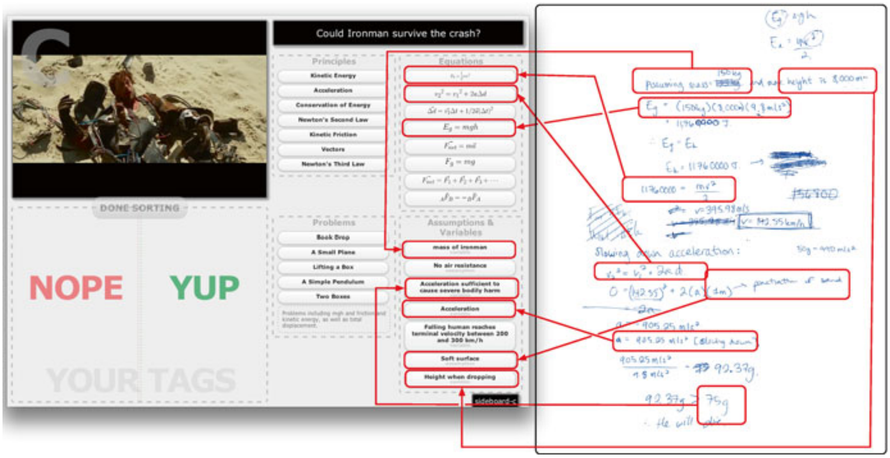

Fig. 12.13 This shows a group's final worksheet for solving their challenge problem. The *red* boxes highlight which elements (i.e., equations, variables, and assumptions) on the worksheet correspond to the codeveloped elements from their zone's interactive display (Color figure online)

process as students noted that "*having the tags and the equations gave [them] a general idea of what the problem related to, so [they] knew the kinds of information to draw from, so it narrowed [their] scope a lot.*"

Evaluating the PLACE Enactment: Did We Support a Knowledge Community Across Contexts?

PLACE was designed as a 12-week curriculum that enabled students to spontaneously and seamlessly connect with an evolving knowledge base across a variety of contexts through carefully scripted interactions. This required that students would be able to contribute and access content when the desire or need arose, but also that such interactions with the knowledge base would be conducted within the course of carefully scripted activities that included various forms of technology scaffolding.

PLACE has been instrumental in supporting our own understandings about how to design and enact such responsive and adaptive curriculum within a well-defined model for learning (KCI). An effective means of evaluating the overall design of PLACE is to examine its ability to achieve its codesigned curricular goals within the context of the KCI model. Thus, we discuss PLACE's enactment in terms of the KCI model below and quote excerpts from both student and teacher exit interviews to support our evaluation.

Within PLACE, students were able to collectively develop the knowledge base through their contributions of examples of physics in their everyday lives (by capturing examples both in their neighborhoods and from the Internet) and to discuss and refine these ideas in parallel. PLACE seamlessly facilitated student engagement with the knowledge base across four very distinct contexts (at home, in their neighborhoods, in class, in the smart classroom), and its design adapted to both students' informational (e.g., providing students with filterable aggregated views of the knowledge base with the Associative Web) and pedagogical (e.g., drawing relevant material from the knowledge base to scaffold student inquiry in the smart classroom) needs.

The evolving knowledge base was not a stand-alone product. Rather, it was used as a resource for both the development of peer challenge problems and to scaffold student inquiry during the culminating smart classroom activity. All of the interactions within PLACE were directly connected to the domain of physics and were further indexed to specific areas of inquiry (the codeveloped principle tags). Student development and the use of the knowledge base were achieved through carefully scripted activities that were sensitive to the context (where the activity took place), the types of interactions (individual, parallel, cooperative, collective) taking place, and a diverse range of media (including laptops, tablets, and interactive large-format displays).

Sarah: *I think that really made us think and made us also realize that there really is physics in everything, because once we got talking with friends to figure out where can I find Newton's First Law, or Second Law, or Third Law, it was really in literally every aspect of our lives… I had the opportunity to talk to students who have analyzed what they can see around them and examples of those laws that they learned in class; talking to them really helped… even working with people who worked on the same [principles] as me, they would have something, examples, that I would have never thought of.*

From the outset, PLACE was designed to address two targeted science learning goals: (1) facilitating students' investigation of science in situations outside of traditional classroom settings, to help them see "science in their everyday lives," and (2) to help students develop a deeper understanding of fourteen "fundamental" principles of their physics curriculum as determined by their teacher. The macro-scripts within PLACE were carefully designed to have students focus on these principles during artifact creation and debate, and PLACE itself has specific prompts and software checks to ensure that these facets of the script were completed by the students. The scripting of student roles (expert categories) and the peer-contributed examples they were expected to review (as part of the in-class review micro-script) ensured that students interacted with a wide cross section of the knowledge base, toward building a comprehensive understanding of the overall domain.

Teacher: *The tagging part of it enables them to share the same language, and I'm quite sure that five years from now if we were to study these kids, they would remember more about Newton's Laws and things like that than a regularly educated kid here at [the school]. I'm kind of sure because they had to tag all those things; those concepts, conservation of energy, and so on would be more in their brains I think – which is kind of neat because* *the sort of stuff I think they should take away with them is the knowledge of those tags and not so much problem solving subbing into equation stuff. But that conceptual learning would be great if kept forever.*

Steph: *[In] PLACE, I remember what I liked personally, and I know this from my own experience, is taking all the things we did and putting them together at the end… where we took all of it in the smart classroom at the end and put it all together, all the different pieces, for me that was the most interesting part.*

In PLACE the teacher's role was clearly specified, within both the macro- and micro-scripts. As described above, within the broader macro-script the teacher was actively engaged in the development of regular homework activities for the students and in monitoring student contributions to the knowledge base toward providing feedback and formal assessment (including grading student contributions and homework using the teacher feedback tools). In the micro-scripting of activities, the teacher was able to look into the class' "state of knowledge" (by reading the student-generated discussion) in order to adjust upcoming class lectures and to engage students in in-class discussion around student-generated artifacts or particular homework questions. During the culminating activity, the teacher used specialized tools to orchestrate the flow of activities in real time.

Teacher: *It was like: Wow I didn't have to explain that before – well that was because I didn't know kids were thinking or confusing that particular thing before.*

Evaluating the S3 Software Agents Within PLACE

A significant outcome of this research was the advancement of the S3 technology infrastructure, which supports knowledge communities across a diversity of contexts and scripted interactions. Central to the ability of S3 to make these interactions possible was the careful design of the intelligent software agents that acted upon the emergent metadata of the knowledge community. This paper advances the notion of three general classes of agents that leverage this metadata toward facilitating both real-time (micro-) and longer duration (macro-) scripted activities: Content Agents, Activity Sequencing Agents, and Grouping Agents. Below we evaluate our implementation of these agents within the PLACE curriculum enactment.

Content Agents

During their creation and debate of physics examples, the personal student tracking agents effectively encouraged students to monitor their own contributions and the growth of ideas of the community. An examination of the server logs and individual students' interactions with artifacts showed that students often did return to their previous contributions after other classmates had acted upon them, indicating a sense of ownership and engagement with their contributions to the knowledge base. The status page was a catalyst for this sustained involvement, by tracking and displaying an individual's contributions and changes to these contributions, PLACE was able to give students a sense of belonging to the community and of the continued growth of ideas.

Pearl: *It was good knowing there was just one place you could go and then finding all your stuff there and just posting your questions. You could see what everyone else was doing; it was easy to evaluate my progress over the year.*

Within the culminating smart classroom activity, agents effectively captured and responded to the emergent metadata of student interactions (i.e., negotiated principle tags and problems), their location within the room (e.g., board A), and other students in the room who shared their location toward distributing script- and context-dependent materials. For example, the *Bucket Agent* distributed materials to facilitate whole group involvement in the task.

Tim: *Well [outside the smart classroom] it would be on a computer screen or something like that, so you'd be like "do this," but we wouldn't be all interacting with it; here I think everyone could all interact with it which was what made a difference.*

Similarly, during the in-class portion of the culminating activity, a *Bucket Agent* used a different pedagogical goal (that of getting through all the material in a time-efficient way), to monitor individual groups and distribute materials across the entire class in a way that allowed the script to be completed within the tight time constraints of a single class period (even with all of the distractions and delays of a "normal" class).

Activity Sequencing Agents

In PLACE, the Activity Sequencing Agents played several major roles in the enactment of the culminating activity. First, the *Consensus Agents* helped students to reach consensus on ideas before moving them to the next task thereby promoting the open discussion of ideas

Sarah: *There was a lot of sharing and applying knowledge, because you had to explain to other people why [a principle or an equation] would apply, and it was kind of recapping your knowledge and also persuading others, expressing your opinion, everything that we did together.*

The *Student Progress Agents* tracked individual, small group, and whole class progression, giving both students and teachers insight into the state of the class within the activity toward reducing orchestrational load. Such agents continuously refreshed the ambient display to show students where they were within the script, when they had completed a phase in the activity, and when the time for an activity had run out. These same agents also alerted the teacher when all the groups had completed a step (on his orchestration tablet), before activating the next step in the script. When the teacher activated the next step on his tablet, the students' individual tablets were instantly updated to reflect the new step in the activity and their own specific roles and locations.

Sarah: *It was a good way to pace everybody and make sure that everyone was going at the same pace.*

Grouping Agents

Within PLACE, the *Grouping Agents* played a central role in the orchestration of the smart classroom activity. The ability for these agents to group students based on specific predefined pedagogical configurations where the students who would fit their conditions could not be known a priori was an exciting outcome of the smart classroom implementation. The grouping and movement of students is a complicated and time-consuming task in any classroom, and being able to not only automate it but to also include processing of emergent patterns (something that would be impossible for a human in real time) provided critical support for managing the class' orchestrational load. In PLACE we successfully demonstrated two such agents (sorting students based on tagging frequency and a modified "jigsaw") which hold promise for more complex ones in future iterations.

Teacher: *It was such a sort of shifting paradigm kind of lesson, with the pacing and, I don't know, just the kinetics and the motion in the room and kids moving around was a lot to follow, [but] I didn't need to worry about it, it was just taken care of by the various technologies.*

Jen: *Well normally the teacher would just say ok and now your next group is this, and they would be the one who would say ok now your time is up and things like that. But with the board it was like ok, this is where we have to go, and that's how much time we have left, so we didn't really need the teacher for that any more… he could just focus more on going around and talking to the groups.*

Transitions Across Contexts: Factors and Design Principles

One outcome of this research is our ability to reflect on how the curricular design supported not only productive interaction *within* different contexts but also the transitions *between* these contexts. Below we describe several design principles for cross-context learning that arose from this intervention. We do not propose that the principles described below are the only possible ones for supporting cross-context learning nor do we suppose that our uses of these principles are the only ones possible; rather given the relatively new domain of this research, we offer our findings as a starting point for other researchers who wish to enact similar designs. We discuss these principles in relation to three transitions that were central to the successful enactment of PLACE, the goals for using materials from previous contexts, the strategy we adopted the script design, and the use of intelligent agents, data structures (e.g., structured and semantic metadata), and data mining.

Visualizations of Community Knowledge

The first transition that we had to consider was between the **individually collected examples** and **collaborative online inquiry activities**. We wanted student-contributed content, rather than materials found in textbooks or other professionally curated materials, to play a meaningful role in the class' inquiry. To this end we designed scripts that specifically required students to draw from the collaboratively constructed knowledge base (e.g., the in-class Challenge Problem Creation script). Our main challenge was finding ways for groups to meaningfully search the large repository of student artifacts to find materials that fit their specific needs. It was in response to this challenge that we build the Associate Web. The associative web was able to mine artifacts from the knowledge base based on their student-assigned tags and present them in a way that was both useful and meaningful for the given context.

We also wanted students at home to see how in-class activities, such as scripted "peer feedback" activities, affected their own contributions to the knowledge base. This was the impetus for the aggregated news feeds, which leveraged system-generated metadata about individual students (e.g., which artifacts they had contributed to or worked on). These different aggregated and filterable views served as a bridge for students to orient themselves within the larger knowledge community when on their own at home.

Design Principle: *Aggregated visualizations of the community knowledge base can play a meaningful role in bridging contexts, but must present the information in ways that are relevant to the context and scripted activity.*

Data Structures and Semantic Metadata Supports

The second transition concerned the movement of materials and student roles between the **at-home stage** (on PLACE.web) and the **in-class stage** (using PLACE.neo tablet apps) of the culminating activity. We needed the small groups to review the work of the individual students and to gain consensus on their assignment of principle tags and equations. In order to do this we needed the system to collect all the individual responses from the at-home stage and aggregate them in ways that allowed students to collaboratively discuss and debate them. Because the underlying metadata was clearly semantically defined (e.g., using metadata structures such as "problems," "principles," "equations"), we were able to easily create views that supported the desired scripted interactions.

These same semantic metadata structures also played a significant role in transitioning the **in-class** artifacts to the **smart classroom**. In the smart classroom, PLACE.neo was able to connect the negotiated tags assigned to a video wall (during step 2, see Fig. 12.8 above), to those attributed to the artifacts from **in-class stage**, to present students with items from the knowledge base that shared the same tags as the video. The ability to leverage the semantic metadata generated in each context allows information to not only move seamlessly between contexts but to also be aggregated in new ways as the knowledge base grows and becomes more interrelated.

Design Principle: *Data structures should be designed to facilitate the organization of student materials for use across different contexts.*

The Orchestrational Role of Intelligent Software Agents

Within the **smart classroom context,** we wanted students to be able to use the materials generated during the **in-class stage** of the script as scaffolds for their problem solving. For this design we knew that the system could not know *a priori* which items would be need by which groups during the activity. In response we needed to develop agents that could draw from the database artifacts that were semantically connected to the students' inquiry and distribute those artifacts evenly to all the group members. Requiring a teacher (or the students) to be aware of every item in the database and their potential connection to the evolving products or real-time inquiry requires a prohibitive level of orchestrational load on participants.

Similarly, during the **in-class stage** we needed to distribute the aggregates of the problems completed during the **at-home stage** (described above). The goal of this activity was not to make sure *every* group saw the same number of problems; rather it was to ensure all the problems we have seen *once* within the confines of a 60-min class. As described above, each time a group gained consensus on one of the aggregated art-home problems, the *Bucket Agent* was able to send another from the set to the group. The ability to quickly assess the state of the activity and draw the required materials from another context in pursuit of the scripted goals provides another layer of adaptive orchestrational support.

This study shows the potential for intelligent software agents to assess complex and changing orchestrational factors such as a student's location (both within and outside the classroom), whether they are working individually or collaboratively; their place within the script; and their past actions to connect them with required materials from the knowledge base. Although our research only engaged these particular agents in two specific spaces (in class and in a smart classroom), the results show promise supporting learners across a wide range of contexts depending on their emergent needs within complex pedagogical scripts.

Design Principle: *Intelligent software agents can help orchestrate class activities that require the retrieval of materials from other contexts based on real-time search conditions or emergent class patterns.*

Conclusions and Future Directions

This study addresses the challenge of developing innovative learning environments for students that blend rich inquiry with the world around them and well-defined pedagogical and curricular goals. We develop technologies that allow students to seamlessly take part in a community whether they are at home, in class, or out playing with friends. How do we transform and aggregate potentially large sets of user-generated data in ways that make sense to students in terms of their progressive knowledge work? How do we script the micro-activities across these contexts to facilitate our longer-tail curricular goals? And what role can intelligent agents play to aid in facilitating the orchestration of these increasingly complex scripts? By designing and enacting PLACE, we have begun to understand the role that agents can play within these systems by providing students with timely insight into their place within the community, suggestions for next steps, the delivery of timely resources, and the grouping and assignment of roles in response to emergent patterns within the class.

As we progress in our investigations, there will be new opportunities for agents to leverage the semantic metadata of the community to create knowledge awareness both for the individual students by more directly connecting them with the relevant products of their peers and the community by producing unforeseen "rise above" trends for further class investigations. These agents have the potential to connect both the long-term investigations of students through persistent portals such as PLACE.web and by making decisions in real-time based on complex student patterns and emergent data that would be impossible to do by hand as in PLACE.neo. As we move forward in these designs, we must be mindful of the role of the teacher within such complex curricula and not relegate them to a role of passive observer or vague instructions to be a "guide on the side." Instead, we must include carefully designed orchestrational supports that empower teachers as active facilitators and role players in the knowledge community.

References

Alavi, H. S., Pierre, D., & Frederic, K. (2009). Distributed awareness for class orchestration. In *Learning in the synergy of multiple disciplines* (pp. 211–225). Berlin/Heidelberg: Springer.

Anderson, T., & Whitelock, D. M. (2004). The educational semantic web: Visioning and practicing the future of education. *Journal of Interactive Media in Education, 2004*(1), 1–15.

Barab, S., & Squire, K. (2004). Design-based research: Putting a stake in the ground. *The Journal of the Learning Sciences, 13*(1), 37–41.

Bielaczyc, K., & Collins, A. (1999). Learning communities in classrooms: A reconceptualization of educational practice. *Instructional-Design Theories and Models: A New Paradigm of Instructional Theory, 2*, 269–292.

Brown, A. L., & Campione, J. C. (1996). *Psychological theory and the design of innovative learning environments: On procedures, principles, and systems*. Mahwah: Lawrence Erlbaum Associates, Inc.

Brusilovsky, P. (2001). Adaptive hypermedia. *User Modeling and User-Adapted Interaction, 11*(1–2), 87–110.

Buckingham, D. (2007). *Beyond technology: Children's learning in the age of digital culture*. Cambridge: Polity Press.

Bull, G., Thompson, A., Searson, M., Garofalo, J., Park, J., Young, C., et al. (2008). Connecting informal and formal learning: Experiences in the age of participatory media. *Contemporary Issues in Technology and Teacher Education, 8*(2), 100–107.

Chan, T. W., Roschelle, J., Hsi, S., Kinshuk, S. M., Brown, T., et al. (2006). One-to-one technology-enhanced learning: An opportunity for global research collaboration. *Enhanced Learning, 01*(01), 3–29.

Chi, M. T., Feltovich, P. J., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science, 5*(2), 121–152.

Cuendet, S., & Dillenbourg, P. (2013). *The benefits and limitations of distributing a tangible interface in a classroom*. Paper presented at the 10th international conference on computer supported collaborative learning, Madison, WI.

Dillenbourg, P., Jarvela, S., & Fischer, F. (2009). The evolution of research on computer-supported collaborative learning. In N. Balacheff, S. Ludvigsen, T. Jong, A. Lazonder, & S. Barnes (Eds.), *Technology-enhanced learning* (pp. 3–19). Dordrecht: Springer.

Dillenbourg, P., Zufferey, G., Alavi, H., Jermann, P., Do-Lenh, S., & Bonnard, Q. (2011). Classroom orchestration: The third circle of usability. *Cscl2011 Proceedings, 1*, 510–517.

Dohn, N. B. (2009). Web 2.0: Inherent tensions and evident challenges for education. *International Journal of Computer-Supported Collaborative Learning, 4*(3), 343–363.

Johnson, R. B., Onwuegbuzie, A. J., & Turner, L. A. (2007). Toward a definition of mixed methods research. *Journal of Mixed Methods Research, 1*(2), 112–133.

Kaplan, F., & Dillenbourg, P. (2010). *Scriptable classrooms. Classroom of the future: Orchestrating collaborative spaces (2005)* (pp. 141–162). Rotterdam: Sense Publishers.

Kohen-Vacs, D., Ronen, M., Ben Aharon, O., & Milrad, M. (2011). Incorporating mobile elements in collaborative pedagogical scripts. In *Proceedings of the International Conference on Computers in Education, ICCE.*

Kuhn, A., McNally, B., Schmoll, S., Cahill, C., Lo, W. T., Quintana, C., & Delen, I. (2012, May). How students find, evaluate and utilize peer-collected annotated multimedia data in science inquiry with Zydeco. In *Proceedings of the 2012 ACM annual conference on human factors in computing systems* (pp. 3061–3070). ACM, New York.

Looi, C., Seow, P., Zhang, B., So, H., Chen, W., & Wong, L. (2010). Leveraging mobile technology for sustainable seamless learning: A research agenda. *British Journal of Educational Technology, 41*(2), 154–169.

Mason, J. (2006). Mixing methods in a qualitatively driven way. *Qualitative Research, 6*(1), 9–25.

Milrad, M., Wong, L.-H., Sharples, M., Hwang, G.-J., Looi, C.-K., & Ogata, H. (2013). Seamless learning: An international perspective on next generation technology enhanced learning. In Z. L. Berge & L. Y. Muilenburg (Eds.), *The handbook of mobile learning* (pp. 95–108). New York: Routledge.

O'Donnell, A. M., & Dansereau, D. F. (1992). Scripted cooperation in student dyads: A method for analyzing and enhancing academic learning and performance. In R. Hertz-Lazarowitz & N. Miller (Eds.), *Interaction in cooperative groups the theoretical anatomy of group learning* (pp. 120–141). Cambridge/New York: Cambridge University Press.

Pea, R. D. (2004). The social and technological dimensions of scaffolding and related theoretical concepts for learning, education, and human activity. *The Journal of the Learning Sciences, 13*(3), 423–451.

Rainie, L. (2012). *Smartphone ownership update: September 2012*. Pew Internet & American Life Project.

Roschelle, J., Penuel, W. R., & Shechtman, N. (2006). *Co-design of innovations with teachers: Definition and dynamics*. In Proceedings of the 7th international conference on learning sciences (pp. 606–612).

Scardamalia, M., & Bereiter, C. (1994). Computer support for knowledge-building communities. *Journal of the Learning Sciences, 3*(3), 265–283.

Scardamalia, M., & Bereiter, C. (2006). Knowledge building: Theory, pedagogy, and technology. In *Cambridge handbook of the learning sciences* (pp. 97–118). New York: Cambridge University Press.

Sefton-Green, J. (2004). *Literature review in informal learning with technology outside school* (Report 7). Bristol: Futurelab.

Sharples, M. (2013). Shared orchestration within and beyond the classroom. *Computers & Education, 69*, 504–506.

Simon, B., Dolog, P., Miklos, Z., Olmedilla, D., & Sintek, M. (2004). Conceptualising smart spaces for learning. *Journal of Interactive in Education, 2004*(1), 1–25.

Slotta, J. D. (2010). Evolving the classrooms of the future: The interplay of pedagogy, technology and community. In K. Makitalo-Siegl, F. Kaplan, J. Zottmann, & F. Fischer (Eds.), *The classroom of the future orchestrating collaborative learning spaces* (pp. 215–242). Rotterdam: Sense Publisher.

Slotta, J. D., & Linn, M. C. (2009). *WISE science: Web-based inquiry in the classroom*. New York/London: Teachers College Press.

Slotta, J. D., & Najafi, H. (2010). Knowledge communities in the classroom. In P. Peterson, E. Baker, & B. McGaw (Eds.), *International encyclopedia of education* (Vol. 8, pp. 189–196). Oxford: Elsevier.

Slotta, J. D., & Najafi, H. (2013). Emerging technologies for the classroom. In C. Mouza & N. Lavigne (Eds.), *Emerging technologies for the classroom: A learning sciences perspective* (pp. 93–112). New York: Springer.

Slotta, J., & Peters, V. (2008, June). A blended model for knowledge communities: Embedding scaffolded inquiry. In *Proceedings of the 8th international conference on international conference for the learning sciences* (Vol. 2, pp. 343–350). International Society of the Learning Sciences, Madison.

Slotta, J. D., Tissenbaum, M., Lui, M., & Zukowski, M. (2012). Smart classrooms for knowledge communities: EPIC technology environment. In *Proceedings of the tenth international conference of the learning sciences*, Sydney (Vol. 2, pp. 64–71). International Society of the Learning Sciences (ISLS), Madison.

Tedjamulia, S. J., Dean, D. L., Olsen, D. R., & Albrecht, C. C. (2005, January). *Motivating content contributions to online communities: Toward a more comprehensive theory*. In HICSS'05. Proceedings of the 38th annual Hawaii international conference on system sciences, 2005. (pp. 193b–193b). IEEE.

Tissenbaum, M., & Slotta, J. D. (2012, July). Scaffolding a knowledge community for high school physics. In *Proceedings of the 10th international conference of the Learning Sciences (ICLS 2012)*, Sydney.

Tissenbaum, M., & Slotta, J. D. (2013). *Researching smart classrooms: Defining technology elements for orchestration of complex inquiry activities*. AERA Annual Meeting, San Francisco.

Tsovaltzi, D., McLaren, B. M., Rummel, N., Scheuer, O., Harrer, A., Pinkwart, N., & Braun, I. (2008, January). Using an adaptive collaboration script to promote conceptual chemistry learning. In *Intelligent tutoring systems* (pp. 709–711). Berlin/Heidelberg: Springer.

Vickery, G., & Wunsch-Vincent, S. (2007). *Participative web and user-created content: Web 2.0 wikis and social networking*. Paris: Organization for Economic Cooperation and Development (OECD).

Wang, F., & Hannafin, M. J. (2005). Technology-enhanced learning environments. *Educational Technology Research & Development, 53*(4), 5–23.

Wheeler, S., Yeomans, P., & Wheeler, D. (2008). The good, the bad and the wiki: Evaluating student-generated content for collaborative learning. *British Journal of Educational Technology, 39*(6), 987–995.

Wiley, D. A. (2000). Connecting learning objects to instructional design theory: A definition, a metaphor, and a taxonomy. *Learning Technology, 2830*(435), 1–35.

Wong, L.-H., & Looi, C.-K. (2011). What seams do we remove in mobile assisted seamless learning? A critical review of the literature. *Computers & Education, 57*(4), 2364–2381.

Zhao, X., & Okamoto, T. (2011). Adaptive multimedia content delivery for context-aware u-learning. *International Journal of Mobile Learning and Organization, 5*(1), 46–63.

Zimmerman, T., & Slotta, J. D. (2003, April 21–25). *Helping students learn difficult biology concepts through knowledge integration activities in the classroom and at an aquarium*. Paper presented to the annual meeting of the American Educational Research Association, Chicago, IL.

Zuboff, S., & Maxim, J. (2004). *The support economy: Why corporations are failing individuals and the next episode of capitalism*. London: Penguin.