Understanding the Power of Augmented Reality for Learning

Gail Carmichael
Carleton University
Ottawa, Canada
gail_c@scs.carleton.ca

Robert Biddle
Carleton University
Ottawa, Canada
robert_biddle@carleton.ca

David Mould
Carleton University
Ottawa, Canada
mould@scs.carleton.ca

Abstract: Augmented reality has recently become a popular interface for various learning applications, but it is not always clear that AR is the right choice. We provide a theoretical grounding that explains the underlying value of AR for learning and identify when it is a suitable interface. Our list of operational design advantages includes AR's use of reality, virtual flexibility, invisible interface, and spatial awareness. This list is backed by four underlying cognitive theories: mental models and distributed, situated, and embodied cognition. We argue that the more design advantages a learning system incorporates, the better AR works as an interface. We also identify a set of questions to be used in the design and evaluation of AR projects. With this, we can begin to design AR for learning more purposefully.

Introduction

Augmented reality (AR) is a type of interface that combines digital objects and information with reality. Virtual content – computer graphics, audio, video, or text – is delivered in a geospatially meaningful way. Many designers have turned to AR for use in applications intended to support learning. From mobile place-based games for science education like Mad City Mystery (Squire K. D., et al., 2007) to Construct3D's (Kaufmann, 2002) collaborative environment with heads-up displays for geometry instruction, there is a large interest in the technology. But what makes AR a good choice? Why is AR well suited for learning?

To answer these questions, we focus on AR's fundamental design advantages and use cognition to identify the core value of AR for learning. We first provide a structured list of the operational advantages of AR. We then examine four cognitive theories and related AR projects to ground these advantages: mental models and distributed, situated, and embodied cognition. Based on this discussion, we determine when AR might be the right type of interface for a particular learning project. All of this leads to a set of concrete design questions that can be used when planning or evaluating an educational AR project.

Related Work

Winn (2003) examines learning with artificial environments, a category that intersects with augmented reality. He suggests that it is worthwhile to take advantage of system theory and neuroscience in order to study learning directly. He offers a framework formed by the concepts of embodiment, embeddedness, and adaptation as an alternative to the more traditional philosophies of learning. Winn recognized the importance of embodiment and embeddedness in artificial environments for learning applications, and we continue to do so in our work.

Shelton and Hedley (2004) consider first person AR that allows for physical manipulation, looking at how the technology supports effective learning of spatial relationships. Spatial cognition theory, animate vision theory, and vision theory are used to argue the advantages of AR in this context. While we look at learning in a more general sense, these concepts connect to our discussion of embodied cognition.

Drawing from experience, Squire et al., (2007) present a set of best practices for designing educational AR games. Several of these are related to our design advantages. For instance, much importance is placed on a connection to real-world locations. Klopfer (2008) also emphasizes place as a key design element that allows for situated and embodied learning with augmented games.

Yuen, Yaoyuneyong, and Johnson (2011) survey the history of AR and its applications, and suggest five main directions AR could take in education. While they discuss general application areas, we focus on *why* AR is well suited for learning.

Billinghurst and Duenser (2012) explore the use of AR in classrooms through example. Though project evaluations have thus far been positive, they posit that future design work needs to be more purposeful:

"Education and sound learning theory will be at the forefront of design decisions from the bottom up." Our work is the first step in fulfilling this call.

Design Advantages of Augmented Reality

According to Azuma et al. (Azuma, et al., 2001), AR "combines real and virtual objects in a real environment; runs interactively, and in real time; and registers (aligns) real and virtual objects with each other." While in context this definition is intended to imply a visual, 3D registration of virtual objects, we adopt it with the broader understanding that location-based alignment of non-visual content is also considered AR. Based on the three properties in the definition, we define four main categories of advantages of AR.

Reality for Free

AR mixes real and virtual objects. As opposed to purely virtual experiences, augmented experiences can be richer and more elaborate because of the deliberate inclusion of real-world objects and behaviors. We can separate this advantage into three distinct elements:

- *Content*. Because the real world is used directly, the AR environment is vast and detailed. A purely virtual experience suffers from a confining bottleneck of content creation.
- *Behavior*. Rich real-world behaviors, such as the laws of physics, are included without effort. In a purely virtual environment, these would need to be simulated by code.
- *Multiple Senses*. Interfaces to purely virtual worlds are typically limited to vision and sound, perhaps with modest haptic feedback; the real world provides a wide array of sensory experiences, including taste, smell, and ambiance such as temperature and humidity.

Virtual Flexibility

The counterpart to the advantages accruing from the use of reality is the flexibility afforded by the use of digital artifacts. Their appearance and behavior are governed by code, and hence can be altered according to the needs of the user or application. The following are two of the most prominent applications of this flexibility:

- *Customization*. Virtual artifacts can be personalized according to user preferences; for example, attributes such as language or color can be adjusted.
- *Impossibility*. Virtual content can depict objects and processes that are impractical or impossible to bring to users otherwise, for instance because they don't exist in the real world, would take too long to observe in real time, or would be dangerous to experience directly.

Invisible Interface

AR users retain their ability to move freely and the interface does not interfere with their ability to observe their real-world environment. They are able to switch attention seamlessly between real and virtual objects; in some cases (such as descriptive markup) the user may not care whether an object is real or virtual. We want to particularly emphasize two aspects of the invisible interface:

- *Natural movement*. User input can take the form of familiar real-world actions; direct manipulation and gesture based interaction are possible.
- *Single focus*. When augmentations are aligned with the task at hand, focus can stay in one place. Users do not, for example, have to change their focus from their task to a paper manual opened beside them.

Spatial Awareness

Some of AR's advantages owe their existence to a clear connection between real and virtual objects, including larger-scale entities such as physical locations. We distinguish two advantages relating to the physical context in which augmentations are embedded:

- *Adjust to surroundings*. Virtual content automatically updates as the user's surroundings change; change can either be a change in the world (such as the movement of a real-world object or a change in temperature) or reflect a change in the user's position or viewpoint.
- *Align spatially*. Close matching of real objects with virtual markup makes associations obvious. For virtual objects which may not have a real analog, opportunistic matching to real objects can lend physicality. For example, labels aligned to a particular object are clearly meant to give information about that object as a real label would.

Figure 1: Example of overlays provided by the Vidente system. Image from the Vidente project site (http://www.vidente.at/), used with permission.

Relationship Between AR's Definition and Design Advantages

Each advantage listed above can be directly connected to one or more characteristic of AR:

- *Combining real and virtual objects in a real environment* allows for the use of real content, behavior, and multiple senses. The fact that virtual objects are involved makes all aspects of virtual flexibility possible.
- *Running interactively and in real-time* means that virtual objects can maintain constant spatial alignment with reality, and the application can automatically react to any changes in the surroundings. Real-time interaction also allows for natural movement.
- *Registering real and virtual objects with each other* leads to the ability to maintain spatial awareness and a single focus.

Cognitive Theories as Theoretical Grounding

To ground the design advantages discussed in the previous section, we examine their relationship to four relevant cognitive theories: mental models and distributed, situated, and embodied cognition. We illustrate the theories with reference to existing AR projects that were either intended for learning or could be used for such a purpose. We also offer a set of design questions that arise from the connections between the cognitive theories and AR; these questions can help guide how an educational AR interface could work.

Mental Models

Broadly speaking, mental models are representations of the world people have learned and use to reason about a particular content domain. Mental models are created based on a person's perception of the world and affect thinking and behavior (Carroll, 2003).

Mental models can be considered *theories*: learners appeal to their existing domain-specific knowledge to determine what to do next, independent of mental mechanisms – i.e., to the contents of their mind. Analogies might be used to apply knowledge of a familiar domain to a new, unfamiliar domain (Carroll, 2003). With AR's *reality for free*, users can appeal to the meaningful relationship of the application to familiar surroundings. *Virtual content* can be designed to match the user's existing mental models. Alternatively, an AR system can use its *virtual flexibility* to alter reality in order to modify an existing theory the user might have.

A magic mirror AR system called mirracle (Tobias Blum, 2012) establishes or corrects a user's mental model of his or her own anatomy. It layers a moveable window over an image of the user's body that shows a 3D model of internal organs or skeleton found within.

Alternatively, mental models are *isomorphisms*, where they closely share the structure of the world they represent. Each object and relationship in the real world has a counterpart in the mental model (Carroll, 2003). AR's *spatial alignment* makes clear the relationships between virtual representations and the real world, allowing the application to externalize the user's mental model. *Virtual flexibility* allows AR to illustrate relationships in an otherwise impossible or impractical situation.

The magic lens GIS system Vidente (Schall, et al., 2008) superimposes geo-spatial information about utilities and related features directly onto images of the ground, allowing workers to see what is underneath and thus build an accurate mental model of the infrastructure (Figure 1). The simulated x-ray vision would be impossible without AR's spatially aligned mixing of real environments and virtual objects.

Also of importance are *mental representations of representational artifacts*: to use a device (the representational artifact), some form of the application domain (where learners' goals reside) must also be represented. Learners need to know much about the device, from its possible operations to how it represents the application domain (Carroll, 2003). In AR, both the device and the application domain can be in the real world that users already understand (*reality for free*), resulting in little or no need to consciously join the two. Alternatively, an application can build users' mental models of the application domain by providing carefully designed virtual objects (*virtual flexibility*) that can be manipulated in a familiar real environment.

The geometry education system Construct3D (Kaufmann, 2002) is designed to help students understand how three-dimensional objects are constructed by guiding them in building shapes. The device space is grounded in the real world, so students have an immediate understanding of how to use the system as they work at building their mental representations in the application domain of constructive geometry.

Design Considerations from Mental Models

- What are your users' existing mental models with respect to what you are trying to teach? How can your virtual design elements best represent the way users understand teachable content according to these mental models? How will their alignment to the real world make their meaning more clear?
- How can you use the virtual flexibility of augmentations to enhance or alter the user's understanding of reality?
- How can the use of the real world make the connection between the interface and the application's problem space easier to understand?

Distributed Cognition

Unlike the theories behind mental models, which focus on cognition at the individual level, distributed cognition (DCog) takes a systems approach. The coordination of all significant features contributing to the accomplishment of a task is important, particularly the surrounding environment, the network of people involved, and the artifacts used (Carroll, 2003).

Representational media in the form of artifacts are distributed units that can contribute to the completion of a learning task. Information can be processed internally, as with mental models, but can also be externalized into physical or virtual artifacts, such as diaries or signs (Hutchins, 1995). Artifacts can be memory or information processing aids to reduce load on a learner's memory or simplify cognitive effort (Carroll, 2003). An interface between an artifact and learner should be designed to match the abilities of that person to the artifact (Norman, 1991). Thanks to augmented reality's *virtual flexibility*, artifacts can be easily customized to suit the user's needs: e.g., language and cultural norms can be respected. AR's automatic adjustment to the user's surroundings (*spatial awareness*) ensures that virtual artifacts will appear when and where they are needed with a *single focus*.

Matsutomo et al. (2012) created an AR application for learning electromagnetics. A visualization of the magnetic field is overlaid on the image of a real magnet that can be manipulated by the user. The visualization is a cognitive artifact, doing the work of figuring out the magnetic field for the learner.

Well-designed artifacts support *activity flow*. Users of a good artifact become unconscious of the artifact itself and imagine themselves as directly operating on the main task. Activity flow can be disrupted, bringing a learner's conscious attention away from the task, potentially impairing results. If the task, artifact, and environment act together cohesively, activity flow can be achieved (Norman, 1991). Augmented reality's *invisible interface* supports *single focus*; this design advantage ultimately allows for an uninterrupted activity flow when using *spatially aligned* virtual artifacts to accomplish a task.

Algorithms developed by Mountney (2010) determine the 3D structure of deforming tissue and project information about surgical procedures into a surgeon's field of view. By providing this information directly on top of the tissue involved in a procedure, surgical operators do not have to break their focus or interrupt their activity flow to retrieve it.

Information and problem solving can also be distributed across a *network of people* wherein divisions of labor can be adapted to the situation, thus reorganizing cognitive artifacts and social context (Carroll, 2003). Augmented reality's *virtual flexibility* allows for the customization of information as it propagates through a group, ensuring a universally understood meaning for both sender and receiver. AR's *spatial alignment* and ability to *adjust to surroundings* ensure that information needed by a person with a particular role is available, even if the person in that role changes.


TASK: Replace Starter Motor

Step 4 of 9: Install new starter motor in starter mount

CAUTION!

Shaft misalignment can damage motor. Ensure holes are aligned with studs in starter mount.

Figure 2: Instructional overlays in the Augmented Reality for Maintenance and Repair (ARMAR) system. Photo used with permission of S. Henderson, Columbia University Graphics and User Interfaces Lab.

Alien Contact! is an educational AR game designed to teach math, language arts, and scientific literacy skills to middle and high school students (Dunleavy, Dede, & Mitchell, 2009). Students play on teams and are assigned the role of chemist, cryptologist, computer hacker, or FBI agent. Tasks and knowledge are distributed accordingly, with the virtual content of the game customized for each person’s role. A network is created that must share and piece together individual virtual and spatially aligned pieces of evidence.

Design Considerations from Distributed Cognition

- What artifacts were involved in the learning task as it was previously performed? Should they be real (to get reality for free) or virtual (to get virtual flexibility) in the AR implementation? Should new artifacts be added?
- How can your application's artifacts use AR's spatial alignment and invisible interface to support focus on the current task? Should your artifacts purposely break activity flow to alert learners to important information?
- Will your application distribute context- and location-dependent information across a network of people? Is it important that learners have access to customized information tied to a specific location or registered to a particular object?

Situated Cognition

Mental models are internal representations, while DCog is concerned with entire systems involved in completing a task. In contrast, situated cognition places importance on where a task takes place, both in terms of the community involved in the activity and the physical world that gives it context.

Taking advantage of situated cognition often means gaining access to a *community of practice*. For example, in many trades, apprentices learn by working with other apprentices, observing those more skilled than themselves, and participating in activities that build up to real tasks (Lave & Wenger, 1991). AR's *invisible interface* and inclusion of *reality for free* allows users to maintain uninhibited access to a community of practice as they work or learn. The AR interface need not inhibit face to face communication, collaboration, or established work practices, due to AR's *spatial alignment* and *single focus*.

A study by Anastassova and Burkhardt (2009) found that automotive repair technicians’ training forms a community of practice as described by Lave and Wenger. A proposal for the design of a training system based on the study suggests replacing paper-based materials with an AR system. One reason for this is to allow for collaboration and information sharing without any hindrance to the technicians’ usual working environment.

Authentic activity is important for effective situated learning. Its characteristics include real-world relevance, multiple complex tasks, the opportunity to collaborate and reflect, and the power to motivate (Reeves, Herrington, & Oliver, 2002). *Reality for free* takes advantage of the availability of the real world to make interaction with the application closer to real-world experience. In particular, the use of *multiple senses* adds a dimension not normally possible with digital interfaces. *Virtual flexibility* means that scenarios that were impossible or impractical to present are now feasible (*impossibility*). Structure can be imposed on top of an unpredictable world and provide hidden information to multiple users, aided by AR's *spatial awareness*.

Many educational AR projects offer some form of authentic activity to engage learners. An example is Mad City Mystery (Squire & Jan, 2007), a murder mystery game that tested whether the inquiry-based tasks introduced in the game would feel authentic to students and help them become part of the community of practice of adult scientific discourse. Students found that the game's meaningful blend of real and virtual increased the game's authenticity.

A different view of situated cognition focuses on the importance of putting *knowledge in context*. Many people can use terms and ideas that connect to their own experiences much more easily than they can think abstractly. Situated learners actively construct knowledge in the context of culture and real-world situations; ultimately, skills gained in context are more transferable to new scenarios (Brown, Collins, & Duguid, 1989; Lave & Wenger, 1991). AR's *reality for free* supports situated cognition by offering real-world context to virtual data. The *spatial alignment* of augmentations to real locations or objects provides easily understood context. For example, terminology can be displayed with the objects described, and interactive explanations of complex systems can be presented directly with those systems.

The Augmented Reality for Maintenance and Repair (ARMAR) system (Henderson & Feiner, 2009) was designed to support armored personnel mechanics in the United States Marine Corps and focuses on the use of explanatory digital artifacts viewed through a head-mounted display. The system overlays the actual machine parts that need attention with a visualization of what to do with those parts, providing suitable context (Figure 2). User studies showed that situating the augmentations in this way allowed the mechanics to work more efficiently.

Design Considerations from Situated Cognition

- Does your target audience learn in an established community of practice? How can you effectively use AR's invisible interface to support the users' activities without getting in the way of that community's communication and work practices?
- How can you use a real-world setting to offer authentic activity to your users? What additional benefits will a spatially aligned digital layer offer?
- How will your users benefit if your application's virtual content is given a clear real-world context? Can you make this content less abstract or otherwise more easily understood by meaningfully aligning it to reality?

Embodied Cognition

The key to embodied cognition is that the human body itself can provide or support models of interaction with the world. A person can perform bodily actions and then repeatedly consult the local environment to make sense of the world (Clark, 1999).

Knowledge in the world implies that intelligence can be found in the interface between the body and the world rather than only in the mind. Traditionally, intelligence is seen as a linear processing cycle wherein a problem is solved in the mind and the body is then instructed on how to act ("perceive, compute, and act"). Embodied cognition suggests instead that problem solving occurs through coordination between the mind and world, maintained by multiple, real-time adjustments (Clark, 1999). AR's *invisible interface* allows users to continue interacting with the world and use the knowledge present in the surroundings. The *spatial awareness* of augmentations allows them to become part of the existing environment. Because augmentations are *aware of user context* and *adjust to their surroundings*, users can continually change their interactions with the world to gain a better understanding of the application and the environment.

The Touring Machine (Feiner, MacIntyre, Höllerer, & Webster, 1997) lets users tour an urban setting and learn more about buildings and landmarks. A heads-up display labels points of interest, and a hand-held computer provides access to relevant information (such as historical facts). The information gets context from being situated, and users are able to interact with and learn about the world by simply walking in it.

A person's *proprioception*, or sense of orientation of body and limbs, should be given special consideration in interface design for virtual environments. Three forms of interaction take advantage of this sense: direct manipulation (handling virtual objects directly with one's hands rather than indirect tools), physical mnemonics (storing virtual objects relative to one's own body), and gestural actions (using one's body sense to facilitate recall of actions) (Mine, Brooks, & Sequin, 1997). Augmented reality is able to support all of these. Direct manipulation can be enhanced by aligning manipulable virtual objects with familiar objects found in the environment (*reality for free*). Augmentations can be *aligned spatially* to personal items such as pens or notebooks allowing for meaningful physical mnemonics. Gestures (*natural movement*) can be used not only to take action but to provide visual cues when working with other users, as reality is not hidden.

The mobile collaborative augmented reality system Studierstube (Szalavári, Schmalstieg, Fuhrmann, & Gervautz, 1998) – used with the aforementioned Construct3D – has a 3D user interface management system that supports the arrangement of applications around users' bodies, creating a physical mnemonic customized to the

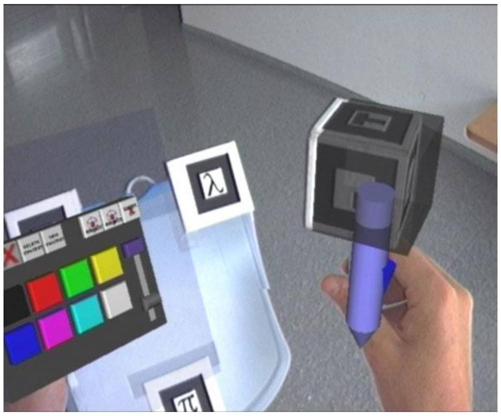

Figure 3: The Studierstube system uses physical mnemonics in its interface. Image from Studierstube project site (http://studierstube.icg.tugraz.at/projects/mobile/), used with permission.

individual user. Applications are positioned by the user and remain anchored there; they are easily found again thanks to spatial memory and proprioception. The user can also interact directly with virtual information via a familiar real-world pen and pad of paper (Figure 3).

Embodied cognition supports an *enactive* approach, where "perception exists in perceptually guided action" (Varela, Thompson, & Rosch, 1992). Active sensing allows a person to gain more information about the surrounding environment, yielding better perceptual data. Epistemic actions can be used as active sensing: exploratory action alters the world in order to get information needed to solve a problem, which may be easier than computing and interpreting information solely in the mind (Anderson, 2003). Many augmented reality applications implicitly make use of the enactive approach, allowing users to alter their perception of virtual data by simply changing their own viewpoint or body position (*align spatially*, *spatial awareness*). AR's *virtual flexibility* can also allow users to actively change the information they see about the surrounding world through augmentations. Users can move in real-world space in a natural and exploratory way to learn more about their surroundings or the data offered by the AR application (*natural movement*).

SandScape and Illuminating Clay (Ishii, et al., 2004) are tangible models used in landscape design. As users shape and mold the clay, the system interprets the resulting geometry and augments it with projections showing water flow and landscape slope values. The systems support an enactive approach as the user physically changes the configuration of the landscape and immediately sees the results. Manipulating the physical model is easier than mentally computing new geographical information.

Design Considerations from Embodied Cognition

- How can you incorporate augmentations into the learner's surroundings such that they become a form of knowledge in the world? Is the learner able to physically interact with these augmentations in real-time in order to solve problems?
- Where is it appropriate to take advantage of proprioception in your interface? Can you support the direct manipulation of real-world objects that have been enhanced by augmentations? What physical mnemonics would be useful? Should augmentations be associated with recognizable items or simply be positioned relative to the body? Is there a good mapping between user actions and gestures they might use to take them?
- What kinds of information can your application provide through epistemic action? Does altering the world yield meaningful changes in the virtual content?

Summary of Cognitive Theories and the Design Advantages

- *Reality for free*. The mental models learners have of the real world make the connection between a device and its application domain better understood. Real artifacts from distributed cognition can be used in any task. The real world can provide situated context to virtual content as well as authentic activity. Users interact unhindered with knowledge in the world through embodied cognition.
- *Virtual flexibility*. Virtual content can be designed to reflect existing mental models or build new ones. Distributed knowledge can be customized or hidden altogether for particular roles in a network of people.

- *Invisible interface*. Distributed cognition's artifacts can be used with a single focus. Learners can work uninhibited within an existing community of practice. Proprioception, enaction, and epistemic action are all supported.
- *Spatial awareness*. Virtual content can be aligned with reality to help build mental models of real world systems. Spatially aligned artifacts appear when and where they are needed, and support activity flow. Spatial awareness of augmentations gives situated context to virtual content, which can become knowledge in the world. Physical mnemonics and epistemic action are also supported.

When is AR the Right Choice?

The cognitive theories discussed above grounded our list of design advantages. These advantages can be used to decide whether augmented reality is a good interface for a particular type of learning scenario. The applications that most benefit from AR are generally those that make natural use of all four categories of advantages.

Virtual flexibility is inherently desirable in all digital tools and applications we design. Similarly, it is desirable to incorporate aspects of the *invisible interface*. We do not want our users to unnecessarily switch their attention between multiple artifacts, for example. We also often try to incorporate natural movements, direct manipulation, and gestures in our interfaces.

In contrast, not all applications need to make use of *spatial awareness* or even have an environment to align virtual objects in. For AR to be a good choice there must be a clear and meaningful relationship between virtual objects and the real world. This might mean that a virtual object is attached to an explicitly related object or location. For example, a digital label that describes the object must be spatially aligned to it to make sense. Or, the relationship may come from seamlessly integrating the virtual object into the environment. For instance, a virtual animal might be shown as though it were present in its real habitat, giving learners the opportunity to observe it. The cognitive theories above offer some suggestions as to where such relationships may be useful, such as when building new mental models or providing situated meaning to the virtual data. But if there is no good reason to associate the virtual objects with some aspect of reality, then AR is likely not best for the application.

Whether an application needs *reality for free* is also an important consideration. It is advantageous to make use of the real world when details found in reality are key to the application. Including reality as it is rather than building a virtual representation of it saves programming effort and reduces the risk that important details about content or behavior are left out. This can be critical when the application involves a task requiring specific behavior with real-world objects. When training for or performing surgery, for example, the exact dynamics, texture, and color of human tissue would be difficult to simulate, yet may be important to the surgeon.

Because *virtual flexibility* and the *invisible interface* represent goals we have for all applications, spatial awareness and reality for free offer the best insight into when to choose augmented reality over other options. If an application can't clearly take advantage of these, then there is likely a more suitable interface type, as is the case with abstract domains in which users do not interact with tangible, real-world objects or data. Instead, users work with virtual information (like data on a computer) or physical abstractions of reality (such as charts). When reality does not play a prominent role in the application, it is difficult to make a meaningful connection between virtual and real objects. For instance, some examples of AR artificially create a connection to reality by having users hold a specially designed card upon which a 3D model will be displayed. While the method of interaction allows an enactive approach to viewing the model, the same could be accomplished with fully digital interfaces that support natural gestures. It is not clear that augmented reality is well used for this kind of application.

On the other hand, applications designed to support learning tasks that are already centered on the real world can make good use of both *reality for free* and *spatial awareness*. Information or problem solving aids can be tied to the relevant aspects of the real environment, as for learning about car engine repair: virtual labels can identify components of the engine while visual instructions can guide the learner. It would not be as easy to do this task virtually given the physical changes made in the real world. The virtual objects have a clear connection to the engine parts in the real world. Many learning applications also benefit from these advantages when, for instance, real world context is important. For example, an application designed to teach photography could use augmentations to illustrate important concepts of composition, depth of field, and so on with visualizations overlaid on the actual scene being photographed. Though these concepts can be illustrated with photographs already taken, interacting with them in real time in the real world makes their context much clearer and helps build a much better mental model. For both these applications, AR is a strong choice.

Design Questions

The more questions a designer can answer yes to here, the stronger the case for using AR becomes.

- Is there a real-world environment that the application or associated task is or should be set in?
- Is there a strong, non-arbitrary association between the virtual data and objects your application uses and some aspect of the environment?
- Is it important that details of the environment, from content to behavior, be preserved?
- If the application supports learning a specific task, is this a non-abstract task that is already performed in the real world?
- Does the application benefit from real-world context?

Conclusion

Augmented reality has been quite effective at capturing the imagination of learners and educators alike, but has not yet been put on a solid theoretical foundation. No one has explained the fundamental reason for AR's value in a learning setting. We have accomplished this with our list of AR's design advantages and our discussion of the cognitive theories that back it. We have identified a set of questions that can be used to evaluate and influence the design of current and future projects. We also addressed the question of when augmented reality is the best choice: learning scenarios that make good use of both *reality for free* and *spatial awareness*, such as those that support concrete, real-world tasks or that clearly benefit from real-world context, tend to best warrant the use of AR.

References

Anastassova, M., & Burkhardt, J.-M. (2009). Automotive technicians' training as a community-of-practice: Implications for the design of an augmented reality teaching aid. *Applied Ergonomics, 40*(4), 713–721.

Anderson, M. (2003). Embodied cognition: A field guide. *Artif. Intell.*, 91-130.

Azuma, R., Baillot, Y., Behringer, R., Feiner, S., Julier, S., & MacIntyre, B. (2001). Recent advances in augmented reality. *Computer Graphics and Applications, IEEE, 21*, 34-47.

Billinghurst, M., & Duenser, A. (2012). Augmented Reality in the Classroom . *Computer*, 56 - 63.

Brown, J. S., Collins, A., & Duguid, P. (1989). Situated Cognition and the Culture of Learning. *Educational Researcher*, 32-42.

Carroll, J. M. (Ed.). (2003). *HCI Models, Theories, and Frameworks: Toward a Multidisciplinary Science (Interactive Technologies).* Morgan Kaufmann.

Clark, A. (1999). An Embodied Cognitive Science? *Trends in Cognitive Sciences*, 345-351.

Dunleavy, M., Dede, C., & Mitchell, R. (2009). Affordances and Limitations of Immersive Participatory Augmented Reality Simulations for Teaching and Learning. *Journal of Science Education and Technology*, 7-22.

Feiner, S., MacIntyre, B., Höllerer, T., & Webster, A. (1997). A touring machine: Prototyping 3D mobile augmented reality systems for exploring the urban environment. *Personal Technologies*, 208-217.

Henderson, S., & Feiner, S. (2009). Evaluating the Benefits of Augmented Reality for Task Localization in Maintenance of an Armored Personnel Carrier Turret. *Proceedings of IEEE International Symposium on Mixed and Augmented Reality (ISMAR '09)*, (pp. 135-144).

Hutchins, E. (1995). *Cognition in the Wild.* MIT Press.

Ishii, H., Ratti, C., Piper, B., Wang, Y., Biderman, A., & Ben-Joseph, E. (2004). Bringing Clay and Sand into Digital Design — Continuous Tangible user Interfaces. *BT Technology Journal*, 287-299.

Kaufmann, H. (2002). Construct3D: an augmented reality application for mathematics and geometry education. *Proceedings of the tenth ACM international conference on Multimedia* (pp. 656-657). ACM.

Klopfer, E. (2008). *Augmented Learning: Research and Design of Mobile Educational Games.* MIT Press.

Lave, J., & Wenger, E. (1991). *Situated Learning: Legitimate Peripheral Participation.* Cambridge University Press.

Matsutomo, S., Miyauchi, T., Noguchi, S., & Yamashita, H. (2012). Real-Time Visualization System of Magnetic Field Utilizing Augmented Reality Technology for Education . *Magnetics, IEEE Transactions on*, 531 - 534.

Mine, M., Brooks, F., & Sequin, C. (1997). Moving objects in space: exploiting proprioception in virtual-environment interaction. *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, (pp. 19-26).

Mountney, P., & Yang, G. (2010). Motion compensated SLAM for image guided surgery. *Med Image Comput Comput Assist Interv.*, 496-504.

Norman, D. A. (1991). Cognitive Artifacts. In *Designing Interaction: Psychology at the Human-Computer Interface* (pp. 17-38). Cambridge University Press.

Reeves, T., Herrington, J., & Oliver, R. (2002). Authentic activities and online learning. *Quality conversations: Research and Development in Higher Education*, 562-567.

Schall, G., Mendez, E., Kruijff, E., Veas, E., Junghanns, S., Reitinger, B., & Schmalstieg, D. (2008). Handheld Augmented Reality for Underground Infrastructure Visualization. *Personal and Ubiquitous Computing, Special Issue on Mobile Spatial Interaction*.

Shelton, B., & Hedley, N. (2004). Exploring a cognitive basis for learning spatial relationships with augmented reality. *Technology, Instruction, Cognition and Learning*, 323-357.

Squire, K. D., Jan, M., Wagler, M., Martin, J., Devane, B., & al., e. (2007). Wherever you go, there you are: Place-based augmented reality games for learning. In E. Shelton, & D. Wiley, *The Design and Use of Simulation Computer Games in Education* (pp. 265-294). Rotterdam: Sense Publishers.

Squire, K., & Jan, M. (2007). Mad City Mystery: Developing Scientific Argumentation Skills with a Place-based Augmented Reality Game on Handheld Computers. *Journal of Science Education and Technology*, 5-29.

Szalavári, Z., Schmalstieg, D., Fuhrmann, A., & Gervautz, M. (1998). "Studierstube": An environment for collaboration in augmented reality. *Virtual Reality*, 37-48.

Tobias Blum, V. K. (2012). mirracle: An augmented reality magic mirror system for anatomy education. *Virtual Reality Conference* (pp. 115-116). IEEE.

Varela, F., Thompson, E., & Rosch, E. (1992). *The Embodied Mind.* MIT Press.

Winn, W. (2003). Learning in artificial environments: Embodiment, embeddedness and dynamic adaptation. *Technology, Instruction, Cognition and Learning*, 87-114.

Yuen, S., Yaoyuneyong, G., & Johnson, E. (2011). Augmented reality: An overview and five directions for AR in education. *Journal of Educational Technology Development and Exchange, 4*(1), 119-14.