A Sketching Game for Art History Instruction

Rajinder Wasson* Carleton University
David Mould† Carleton University
Robert Biddle‡ Carleton University
Cristina S. Martinez§ *Independent Scholar*

Abstract

Study of images, in context as widespread as medicine, geography, and art history, demands attention to detail and exercise of memory. Presented with traditional textbooks, students have difficulty concentrating for the long periods of time needed to absorb the information. This paper presents a sketch-based game for learning images. We tested the game in the domain of art history, and in a controlled user study, we found that it dramatically increased the time that participants chose to spend learning about historical paintings. The increase in study time was accompanied by a proportionate increase in information retention and familiarity with the images.

CR Categories: H.5.1. [Multimedia Information Systems]: User Interfaces—;

Keywords: sketch-based interfaces; design; games; gamification; art history

1 Introduction

Introductory courses on art history may introduce students to hundreds of different artistic images, and students are expected to learn a significant body of knowledge about the set of images encountered in the course. Each image has many points of interest and associated information; students should not only be able to recall specific details, but should gain a general familiarity and fluency with the image collection. This can only be done by dedicating long hours of study. Yet students have difficulty sustaining their effort to attain sufficient mastery of the subject matter. This is not to say that it is impossible to do so, merely that in practice few students do so to the satisfaction of their instructors.

In contrast, computer games have an incredible power to persuade their players to pursue endless repetitive tasks. Tales of the countless hours spent playing solitaire or *World of Warcraft* are legendary. Games can draw players to continue despite an overwhelming repetitiveness in the tasks being performed. Ian Bogost's deliberately farcical *Cow Clicker* [Tanz 2011] is an extreme example, where players click monotonously on the picture of a cow, eventually earning enough game currency to obtain a different cow, which can then be clicked on in its turn. While many commercial and educational games have elaborate activities and puzzles designed to captivate players, *Cow Clicker* and other examples suggest that games can in fact be quite simple while retaining their power. In our case, players are expected to have some initial desire to learn the content already; we reasoned that we require only a very thin veneer of game-like activities to hold players' attention while they examine and learn the art history images. This paper discusses our design for an educational art history game. We drew on the expertise of one author in particular, who is an art history instructor and who provided the domain knowledge for the game.

Our design rests on three pillars. First is a mechanic called *prospecting*, a repeated epistemic action where players search for hidden information in an in-game landscape. Second is the notion of *purification*, a term coined by Jesse Schell [Schell 2008] to describe a compulsion to complete a task when that task is divided into small pieces with visible representations. Finally, the activity of *sketching*, where players sketch continous paths along predefined contours within the image plane, following possibly meaningful shapes called *glyphs*. We conducted a 25-participant user study in which we compared our design with a digital version of a standard textbook view, where the image and all associated lessons were presented at once. We hoped that participants would spend longer on images seen in the game view and that there would not be a severe reduction in learning rate in the game condition. These hopes were borne out: participants chose to spend more than three times as long on the game images, and we saw no evidence of a difference in learning rate between the two conditions.

The contribution of this work is a design for software for learning images, tested in the domain of art history but believed applicable to any domain in which information is or can be presented primarily in the form of 2D images. In the remainder of the paper, we discuss our design in detail, describe our user study and its results, and identify particular strengths and weaknesses of the approach we took.

2 Previous Work

The textbook and lecture continue to be central to art history courses [Sowell 2009]. Availability of high-resolution images and video, accessible, for example, from the Google Art Project [Google 2011] or the website of the Louvre [Louvre 2013], provide new opportunities for study of the source material. However, anecdotally, motivation continues to be a problem. Students have difficulty finding the desire to spend long periods of time studying, and even when large blocks of time are available, concentration can be spotty. Yet when hundreds of images must be learned, as is common in an introductory art history course, there is no alternative to spending time.

Sketching is often linked with exploratory design or externalization of memory [Bertel et al. 2006]. However, we use it here in an effort to promote more active engagement with preexisting material. Sketching to better study the artwork is a common practice among art historians. Fish and Scrivener [Fish and Scrivener 1990] suggest that sketching is linked to memory: the physical interaction focusses attention and can make it easier to remember details. We hope to take advantage of this effect by requiring users to sketch out shapes in the images to be studied. In general, we argue that intentional action related to an object of study leads to more complete recall than mere observation. For example, anecdotally, singing along to a song in *Rock Band* [Electronic Arts 2008] leads to much faster mastery of the lyrics than repeated passive listening. Paulson et al. [Paulson et al. 2008] created a suite of educational

*e-mail:raj.wasson@gmail.com
†e-mail:mould@scs.carleton.ca
‡e-mail:robert.biddle@carleton.ca
§e-mail:martinezcsm@gmail.com

activities with sketching interfaces, where the system would provide automated feedback to students' responses to questions. Piper and Hollan [Piper and Hollan 2009] used digital sketching to motivate and encourage undergraduate students in studying cognitive science. Even unrelated sketching-type actions may assist memory: Andrade [Andrade 2010] reports that doodling as a secondary task improves recall in a primary listening task, possibly by occupying excess cognitive capacity without competing with language processing.

Existing games for learning art history typically lack interaction with the images: they are generally either quiz-type games that test knowledge but do little to impart it, or adventure games which seek to teach incidentally in the course of following the game's story. A typical example of the latter is Eduweb's *A. Pintura: Art Detective* [Eduweb 1997], in which the player is tasked with investigating a series of mysteries, solved by comparing the composition, colors, and subject matter of different paintings.

Egenfeldt-Nielsen [Egenfeldt-Nielsen 2011] highlights three critical factors for a successful educational game: *integration*, where gameplay and intended learning outcomes are tightly coupled; *motivation*, achieved by providing quick feedback and a close connection between action and reward; and *focus*, an unrelenting emphasis on the content to be learned. For our game, we sought to include these elements. We also drew on Schell's notion of *purification* [Schell 2008] to motivate the player to continue. Purification refers to the process of completing a finite set of requirements, such as destroying all the enemies in *Space Invaders* or clearing all the dots in *Pac-Man*. When presented with distinct tokens as in these examples, the player will feel motivated to complete the task by gathering the last dots or clearing the remnants of the enemy fleet. Our own design relies in part on purification, as we discuss in the next section.

3 Design

We created a software application for studying art history, intended to encourage users to spend more time looking at the images. Users can choose an image (a digital version of an artwork) from among those available. Users then search through the image for invisible glyphs: glyphs are arbitrary marks we inserted into the image, usually associated with a significant image element. A glyph is revealed when the user clicks nearby. Once revealed, a glyph can be sketched with the mouse pointer, and successful completion of the sketching is tracked by the software. The user's progress through the image is shown by a set of stars at the bottom of the image.

The search for glyphs is called "prospecting". Users scan the image for likely locations for a glyph; we usually attached the glyphs to visible features, and users will swiftly realize this if they do not guess it immediately. A glyph is found when the user clicks the mouse near the glyph's location. Whenever the user clicks, further clicks are disallowed for a brief interval; after some experimentation, we settled on a delay of five seconds. This is short enough that the user does not abandon the task, but long enough to have an impact. The user does not want to needlessly incur the delay, so will think a bit about where to spend the click. Prospecting, in our usage, is a particular form of exploratory epistemic action, analogous to the action of opening a box in *Minesweeper*.

Our implementation of the prospecting mechanic encourages longer study in two ways. First, it takes some time simply to find all the glyphs, so the user is faced with the image for a certain period of time: the user cannot rush to complete the image even if so inclined. Second, the brief clicking lockdown encourages strategic placement of clicks: users will not click arbitrarily, but will guess glyph locations based on image content. Thus, users are lightly incentivized to understand the image even for their most short-term actions. During prospecting, users gain a broad but shallow impression of the image as a whole.

Each glyph is associated with a "lesson" about the image: a piece of information about the image, the artist, or some contextual or historical information that would enhance the viewer's understanding of the piece. Lessons are expressed in one or two sentences of text and the associated lesson is displayed while the user is sketching its glyph. Lessons for completed glyphs may also be reviewed at the user's convenience. One of the authors is an art history instructor. This individual created 7-10 lessons for each image, with the number varying depending on how much information about that image a typical class would attempt to convey.

During lesson creation, designers draw and place glyphs by hand; whenever possible, we attached them to meaningful image elements. The glyph's lesson can refer to the image element that the glyph surrounds, when the lesson has a concrete referent: for example, the artist may have depicted a specific religious or historical figure, named in the lesson. For example, one lesson for Caravaggio's *The Calling of St. Matthew* identifies St. Matthew himself, commenting that *"St. Matthew appears surprised by the intrusion and responds by pointing to himself. His right hand remains on the coin he had been counting."* However, the glyph's location may not relate to the lesson content in the case that the lesson contains general comments about the image, the artist, or the temporal context of the art, for example. A lesson of this type, also relating to *The Calling of St. Matthew*, states that *"the painting alludes to the Flemish tradition of paintings of money-lenders."* Roughly speaking, we have three types of glyphs: those attached to a particular object in the image plane; those whose shape draws attention to an aspect of the composition; and those whose shape is not connected to the underlying artwork, serving only to form a sketching task at some difficulty level or to be a recognizable symbol.

Even glyphs lacking a connection to the image location provide some learning assistance. The incidental details of the image supply a visual mnemonic for the learner, and the glyph's location and shape provide a spatial connection, enlisting the learner's spatial memory. In Figure 1, we depict some sample glyphs for the painting *The Ambassadors* by Hans Holbein the Younger. In this image, we have distributed the glyphs throughout the image plane; this is desirable both because it forces users to search most of the image during prospecting, and because the user will inspect a variety of regions of the image in detail during sketching. During sketching the user will also study the lesson associated with the current glyph.

We do not only want the users to memorize the text of the lessons, however. Imparting the content of the lessons is a partial goal, but our broader objective is for the students to become deeply familiar with the artworks, and ideally themselves to notice details, aspects of composition, and artistic styles rather than merely parroting the information handed to them. At a minimum, the student should be able to recognize the image later. For all this to be possible, the students must have spent time looking at the images.

The prospecting mechanic encourages longer study, as does our implementation of sketching. Once a glyph is found, the lesson is displayed and the user must sketch out the glyph path to complete it. The sketching must occur at a measured pace: sketching either too quickly or too slowly causes the action to fail, in which case the user must start over. Departing from the path also triggers a failure. Our intention in creating this sketching activity was to require users to spend time contemplating the image, while providing them a simple task to prevent restlessness. By design, the task is not too difficult so they have attention left over. Adept readers can

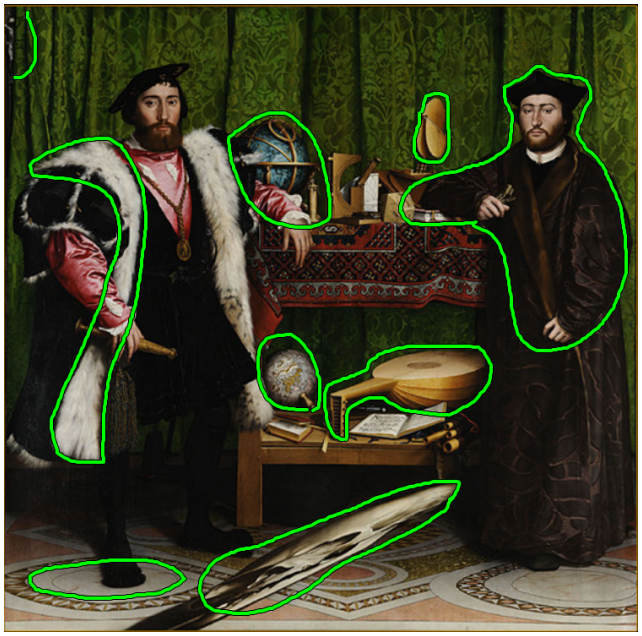

Figure 1: *Sample glyphs in the painting* The Ambassadors. *Glyphs draw attention to specific individuals and objects in the scene, are opportunistically linked with visible structures, or highlight specific features such as the anamorphic skull in the lower middle of the image.*

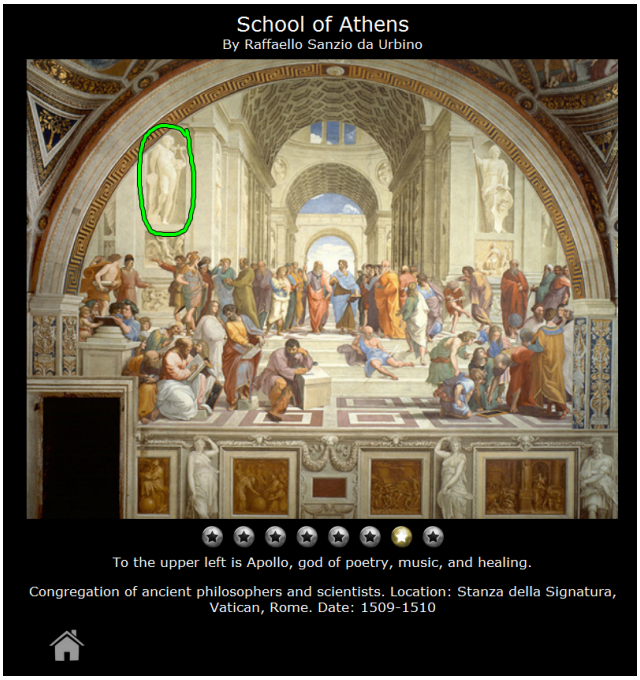

Figure 2: *Screenshot of the game. The player has completed one glyph and is now reviewing the associated lesson.*

take in the lesson in a few seconds, so can be thinking about the lesson while undertaking the sketching. When the glyph is successfully completed, a "hurray" animation plays over the glyph, and the star corresponding to the glyph lights up. Prospecting gave users a broad overview of the image; sketching the glyphs gives them a detailed view of a specific image element and also provides them time to contemplate the associated lesson.

The stars below the image indicate which glyphs have been completed. A star lights up for each completed glyph, so the player can see at a glance whether there is more to do. The lit stars also provide the player access to the glyph and lesson, in case review is desired. The stars are a type of progress bar, but because the glyphs can be completed in arbitrary order, spurious patterns emerge from the sequence of completed glyphs. A basic progress bar provides a mild compulsion to complete the task, but here there is an extra cognitive effect owing to the gaps emerging from semi-random sequences of lit stars. The gaps both provide opportunities for additional purification, when a glyph is completed that links up two formerly separated groups, and add a heightened pressure to continue, so as to fill in an observed gap.

A sample screenshot from the application is shown in Figure 2; some of the main elements of the design are visible here. The majority of the screen is devoted to the artwork itself. The name of the artwork and the artist's name are shown at the top. At the bottom, the user can also see a small piece of explanatory material about the painting: a very brief description of the content, a mention of the physical location of the painting (such as what museum or collection holds it) and the date at which the piece was originally painted. The user's progress through the glyphs is shown immediately beneath the painting as a series of lit and unlit stars; one glyph has so far been completed, and the player is presently reviewing it, so the glyph itself is shown as well as its associated lesson. This lesson refers to an element of the image, so the glyph's shape and location are highly relevant. In the lower left corner, we see a "home base" icon, which can return the player to the hub where other images can be selected for study.

We expected the elements of our design to act in concert to encourage more study of the art history images. Prospecting gives students a broad view of the image. Sketching gives them a detailed view of the sketched part and occupies them with a simple task so that they can attend to the lesson without becoming bored. Purification urges them to continue until all glyphs have been discovered and sketched. To evaluate the hypothesis that students will study more using the game, and to quantify the magnitude of the effect, we conducted a user study, described next.

4 Study

We conducted a user study to determine whether the game and its sketching activity would persuade students to spend more time studying the images, as compared with a textbook-style static visualization of the content. The study was conceived as a within-subjects experiment, where all users saw several images in each of the two conditions, textbook and sketching game. The dependent variables are the times users spent on each image and the number of lessons they recalled about the images afterwards. We collected timing information, recorded user actions such as mouse movement, and used an eye tracker to record gaze location over time. After the subjects had completed a set of images, we asked them to recall the lessons and to make further observations about the images, capturing their responses with audio recordings for later analysis. The timing details and interview responses provide quantitative measures; we used the eye-gaze recordings to get a qualitative measure of image coverage. Full details of the study and our analysis of the results appear in the subsections below.

We marginally altered the standalone game application to fit the needs of our study, instrumenting it to record user activity, making it fullscreen, and slightly adapting the layout. We also built

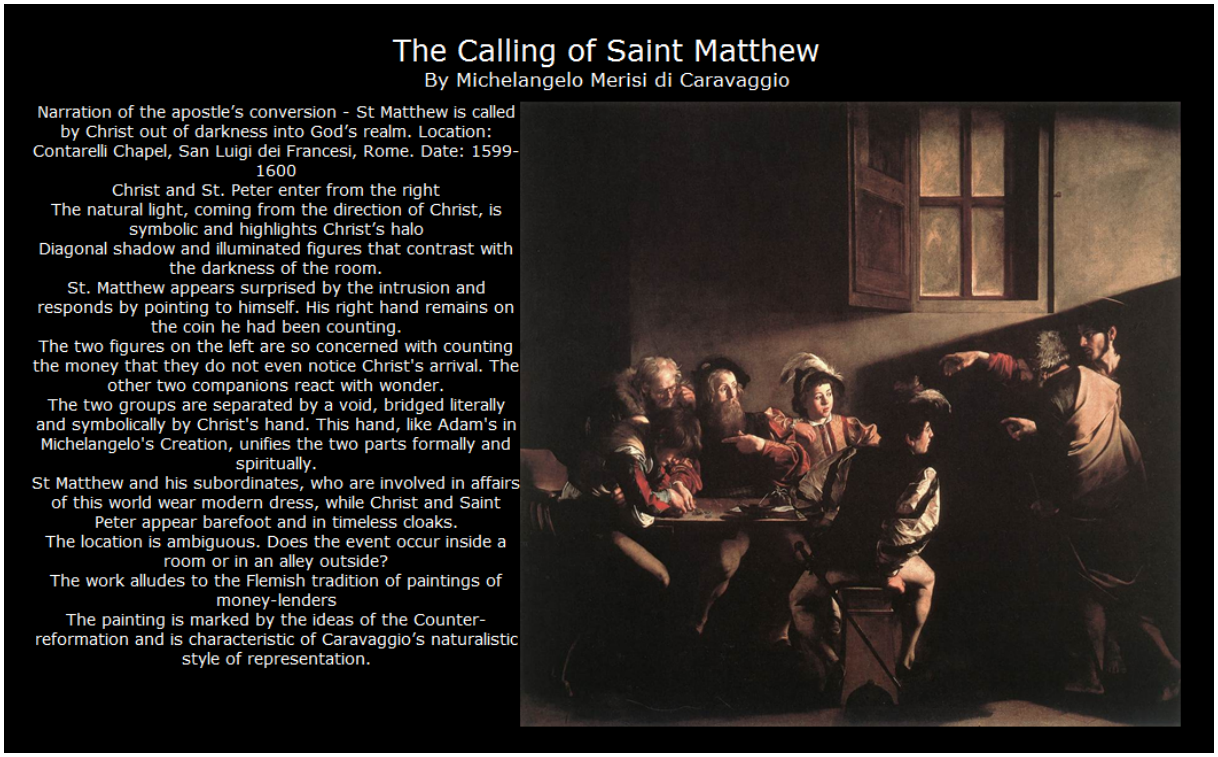

Figure 3: *Screenshot of textbook condition.*

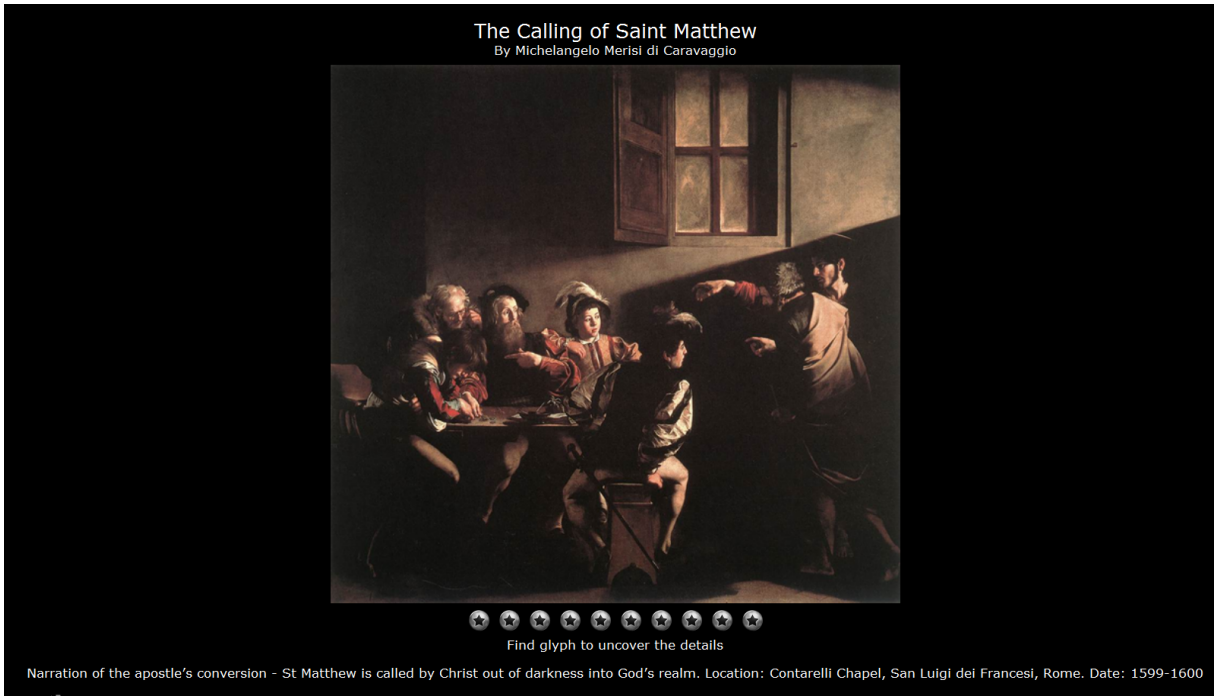

Figure 4: *Screenshot of game condition.*

a "digital textbook" capable of displaying the images and associated lessons directly, without need for prospecting or sketching. In this section we provide details of the study comparing these two methods of information delivery. Our study was approved by the Research Ethics Board of Carleton University.

4.1 Participants

Participants were university students aged 18-25. Art history students were recruited by announcing the study in an introductory art history class, and others were recruited through flyers and word of mouth. Participants were paid $10 for their time. Of the participants included in our analysis, 14 were male and 9 female. Most participants had some regular video game habits (a few hours per week) but only one habitually played more than seven hours per week. This level of involvement with games is typical of the casual game players our design was intended to reach. We recruited 25 participants, but two participants were ultimately excluded, one because of severe lack of sleep affecting performance, and one because of detailed prior knowledge of one of our test images. In a post-test questionnaire, five of the participants acknowledged having previously seen many of the study images, either viewed in class or encountered in television or other media, but did not know much about them, and hence their data was retained.

Approximately half the students were taking or had recently completed an introductory art history class at the time they participated in our study. The remaining half generally lacked any significant prior knowledge of art history. We refer to these two groups as "art history" participants and "other" participants. We included and distinguished these two groups to see whether art history students would perform differently than a general population, perhaps due to differences in interest level and motivation.

4.2 Task

Participants were first given a short questionnaire with questions pertaining to demographics, video game playing habits, and their knowledge of art history. Participants were then introduced to the software and told about the two conditions. They were asked to try to learn all the lessons and informed that they would be tested at the end. Then, they had the opportunity to practice prospecting and sketching on a famous image, Leonardo da Vinci's *Last Supper*. They also did some short exercises to calibrate the eye tracker. Finally, they were given the experimental task: to study all the images in the hub and alert the coordinator when the task was complete.

The image set consisted of 11 images, chosen by the art history instructor (one of the paper's authors) who also provided the glyph lessons and helped determine the glyph shapes and locations. Each participant would see 5 images in the game condition and 5 in the textbook condition; which images were seen in which condition was random on a per-participant basis. Three of the images – two from the ten seen during the body of the task, and a third that had been held in reserve – would be shown at the end and the participants asked to recount their recollections and observations. The two images seen during the exercise were constrained to be in different conditions, ensuring us a supply of data to compare the game and textbook conditions.

The experimental hub (shown in Figure 5) gives a preview of all images and provided summary information about viewing and glyph completion. Subjects would select an image, study it for the desired length of time, then return to the hub and select a new image. They were free to return to images they had previously viewed, though none of the subjects in our study exercised this ability. The participants were able to access the images in arbitrary order; the typical

Figure 5: *Experimental hub: access to images and information about completion. Actual painting images have been replaced by pictures from Flickr.*

strategy, used by nearly all, was to proceed in a systematic way by choosing a starting point and then working clockwise or counterclockwise around the ring. Subjects would not know until first viewing an image whether the image would be studied using the textbook condition or the game condition.

Once participants had completed all images, they were shown three images in turn: one image they had seen in the game condition, one in the textbook condition, and one they had not seen in the study. The same three images were used at this stage for all participants. Participants were prompted to recount all the lessons they could recall about the image plus to make any comments or observations they had. Their responses were audio recorded.

Finally, participants were given a post-experiment questionnaire which elicited their subjective reactions to the game. We will discuss results from this below. Once they had completed the questionnaire, participants were thanked for their time and paid the promised amount.

4.3 Experimental Setup

Participants ran the software in a quiet room free from distractions. The computer was equipped with a 17-inch monitor and a Tobii 1750 eye tracker, also running the Tobii 2.6.0 eye tracking software. Images were portrayed in a window of size 1024x768 (fullscreen); individual images would be presented at a smaller size and with the native aspect ratio of the image file provided by the art history instructor. Our software logged data about each session, including time spent on each image and number of glyphs completed.

4.4 Data

Timing data – tracking how long participants spent viewing particular images – was obtained directly from the automatic logs. We also scored the participants' verbal discussions of the images at the end by reviewing the audio recordings. Participants were given two scores per image: a count of "lessons", i.e., the glyph text recalled in whole or in part, and a count of "observations", which are general statements about details noticed from the images. We were reasonably generous in scoring lessons; even partial recollection was granted a point, although we took care not to double-count any glyphs. We were less generous with observations: observations had to be fairly distinct to be counted, and sometimes a few similar observations were grouped into a single point. Observations might focus on specific details of the painting, such as pointing out a champagne bottle on a table in the scene, or refer to the image as

a whole, such as pointing out that the artist favored warmer colors over cooler colors. We scored the audio without reference to which condition the image had fallen into, merely using the participant ID numbers to match the data back to its condition once the scoring had been completed.

A final source of data was our exit questionnaire, which sought to quantify the participants' subjective reactions to the game. We asked them whether they enjoyed playing the game, whether they thought it would be a plausible means of studying, and allowed them a freeform space to make any comments they thought relevant. (Few availed of this last opportunity, however.) We also inquired whether the participants had seen any of the study images before and attempted to elicit how detailed their previous knowledge of these images might have been. We next discuss the analysis of the data we gathered.

5 Analysis and Results

Our study had 25 participants. Of these, one had detailed prior knowledge of one of the final test images, and one confessed to being impaired by lack of sleep. We excluded these and our analysis pertains to the remaining 23.

The basic question we wanted to answer was whether the participants spent longer on images they saw in the game style compared with those they saw in the textbook style. The mean time per image across all participants in the textbook condition was 1.42 minutes, and the mean time per image in the game condition was 4.81 minutes. This finding is significant ($t(22) = 4.32, p < 0.001$). The participants spent more than three times as long studying the game images as they did the textbook images.

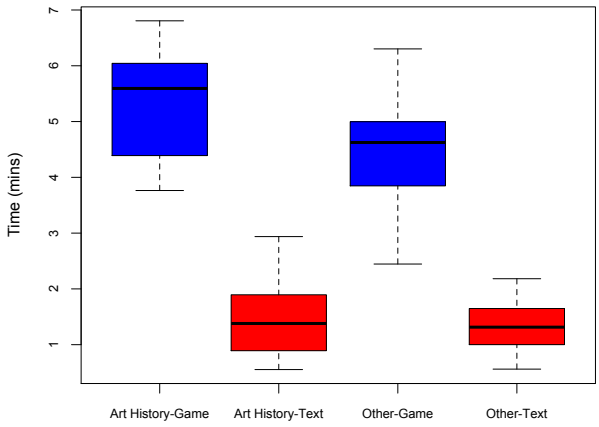

Figure 6: *Box plot of average time spent by participants in each condition.*

We also wanted to see how the art history participants compared to the other participants. Participants from art history spent more time on average than other participants in both conditions. However, this difference averaged only 15 seconds in the textbook condition. In the game condition, the average time spent by art history students was about a minute longer than the others. In neither case was the time difference significant, owing to the wide variation among individuals. A box plot of the times spent appears in Figure 6.

It seems clear that the game condition induced the participants to spend more time. However, this would be rather pointless unless accompanied by more learning: had they learned nothing, we could have had the participants pass the time with any existing entertainments. We therefore compared the lesson and observation counts in the two conditions.

The mean number of lessons recalled in the game condition was 4.83, but the same score for the textbook condition was only 2.74. From a paired t-test, the populations are distinct ($t(22) = 17.06, p < 0.001$). Further, we have the participants' observations, which are more numerous; the mean number of observations in the game condition was 14.2 and in the textbook condition 8.4. This difference is significant ($t(22) = 6.80, p < 0.001$). Observations are important because they show increased comfort and fluency with the images; as mentioned earlier, our goal is not merely to impart specific lessons, but to familiarize the students with the image through prolonged exposure.

Data from the third test image, not previously seen by participants during the study, we had intended to use to benchmark the talkativeness of each participant. The observations made on this third image were generally fewer than for either of the images they had previously seen, and, unsurprisingly, almost no participant mentioned any of the specific lessons prepared for this image (never shown to any participant). There was no discernible pattern in the observation counts for the third image.

There is a positive correlation between the time spent and the amount learned. A scatter plot of lessons recalled vs. time spent shows the relationship: see Figure 7. The same relationship holds for observations vs. time spent, with a similar scatter plot shown in Figure 8. In both plots, we included a least-squares regression line. The qualitative outcome is quite clear: more time spent studying, either in the game or the textbook condition, resulted in greater mastery of the material. The main effect of the game was to convince players to spend more time, hence pushing them further up the curve.

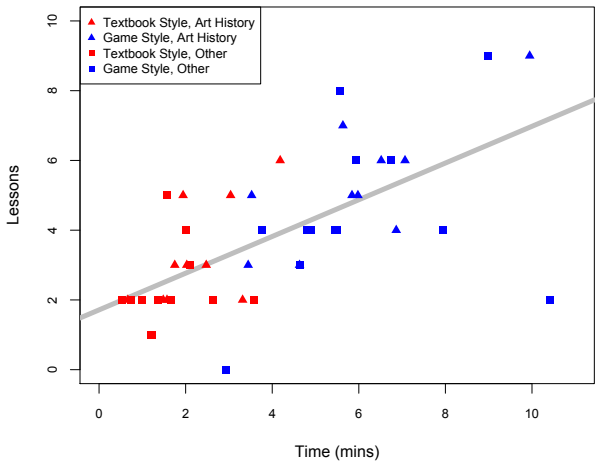

Figure 7: *Scatter plot of lessons remembered vs time spent on image.*

Next we discuss the data recorded by the eye tracker. We provide a qualitative discussion of these results, based on a representative image, *The Calling of St. Matthew* by Caravaggio. *The Calling of*

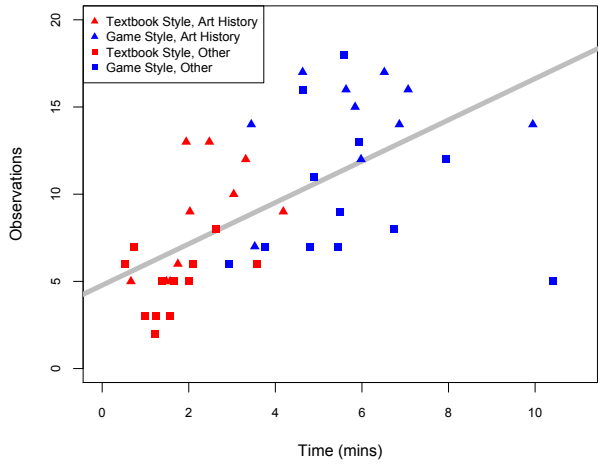

Figure 8: *Scatter plot of observations made vs time spent on image.*

St. Matthew also serves as a test case for our hypothesis that distributing the glyphs throughout the image would cause participants to distribute their attention while searching: as with many other paintings, the most salient parts of the image are near the centre, and the edges less strongly draw the eye. Figure 9 shows the heat map of gaze fixations for the image in the textbook condition and the game condition, overlaid on a screenshot from the condition. In the texbook condition, the image never changes, but in the game condition, we see the player in action, having partially traced one of the glyphs and with the glyph lesson visible in a callout overlaid on the image.

In both conditions, participants' attention was drawn to the figures and faces in the painting. However, the textbook condition mostly lack fixations elsewhere. Conversely, in the game condition, the participants looked at the window, the feet of the figures, and even at areas of the image that at first glance lack any detail at all. The game condition has been successful in getting the participants to move beyond the most obvious image features and to look at other image elements as well, and at the same time, to study the obviously salient features more thoroughly. In the textbook condition, participants mainly just read the text, with seemingly perfunctory visual exploration of the image itself.

In a post-test questionnaire, an overwhelming majority of participants agreed that they enjoyed playing the game (91%). A surprising 43% of participants claimed to be willing to play it on their own time simply as a pastime; 87% of participants said that they would play it in order to learn the material for a class. Approximately two-thirds of our participants (63%) said that they would prefer learning from the game to learning from a textbook. Note that we do not suggest replacing textbooks, rather supplementing them with educational software, so there would be alternatives for students with different preferences. Finally, participants saw the value in the game's central sketching mechanic: nearly all (91%) agreed that the sketching task increased their awareness of specific areas of the image.

We do not consider these subjective responses to carry very much weight. They chiefly serve as a sanity test of the plausibility of our approach. Had the participants held a strongly negative opinion of the game, we should reconsider our design. Their positive opinion frees us to concentrate on the quantitative outcomes.

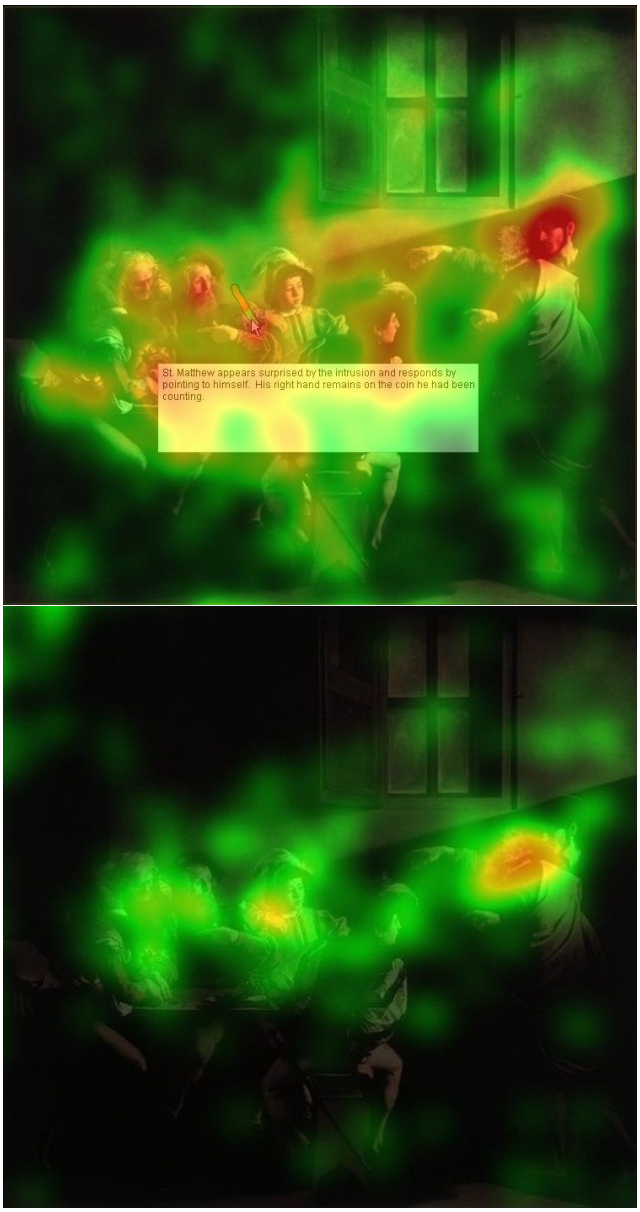

Figure 9: *Attention heatmaps for* The Calling of St, Matthew. *Above: game condition; below: textbook condition.*

6 Discussion

Our design has some major positive aspects. From our point of view, its most important advantage is that it succeeded in its objective: participants did spend more time studying the images – more than three times as long, compared to a simple "digital textbook" presentation of the same information – and the extra time was accompanied by additional learning. In short, participants studied harder when looking at the game than when looking at the static textbook view.

Further, participants stated that they liked the game and would play it in order to learn about the subject matter. Most admitted that they would not play it for entertainment, but our goal is to make studying more palatable, not to make studying preferable to existing entertainments.

There are some limitations, however. The design is aimed strictly at engagement, not necessarily more efficient learning; though we have little quantitative evidence to this effect, we believe that someone could learn more quickly from a textbook than from the game if suitably motivated. The game helps provide motivation to students to continue when they might otherwise abandon their study. We speculate that students might find it easier to learn the material when lessons are attached to locations in the image, since the sketching might engage the students' spatial memory, but this has not been tested. We anticipate a stronger effect from sketching in domains where the marks to be sketched are more critical; an example from geography might be learning states' borders, which would be sketched directly by the students.

The sketching game requires technology (a computer and screen) and is subject to all the disadvantages that entails. A book is more portable, robust, and energy-efficient. Books also enjoy an enormous incumbency advantage: there are huge numbers of existing books on an immense variety of subjects, but content creation is a substantial bottleneck for would-be adopters of our sketching game. Glyphs must be drawn by hand, and, at present, added manually to our database. Our implementation is also somewhat rigid with respect to what information may be associated with a glyph: only short pieces of text are permitted. This is in part a minor limitation of the implementation and in part a consequence of the approach: with limited screen real estate, we can only feasibly display a small image or short text piece alongside the glyph. In principle, however, completing a glyph could unlock access to an arbitrary data store.

The game will not work miracles. It is quite possible for someone to play the game and emerge with little understanding despite the time spent. Indeed, in figures 7 and 8 we can see that the participant who spent the longest time playing the game learned very little. Some inclination to attend the subject matter is still required before any learning can occur.

Lastly, although we strove to find a balance between a sketching task that was not too easy (and hence stave off boredom) and not too hard (allowing participants some leftover attention with which they could attend the image or the lesson), we may have erred on the side of making the task too difficult. Some participants had to retry individual glyphs more than a dozen times, which was not what we had planned. Adaptive difficulty or simply fine-tuning of the static difficulty would be helpful.

We have some suggestions for an extended design, suitable for longer-term deployment but of little use in a short session such as that examined in our study. We can augment the existing purification with an achievement mechanic, providing achievements for completing specific combinations of images. We suggest having multiple levels within the same image, to be unlocked after completing tasks on other images, to have players return to a previously viewed image to reinforce their learning. In a large-scale, long-term use contexts, persuasive mechanics such as leaderboards can be deployed, allowing students to compete with their classmates. A possible partial solution to the content creation issue mentioned above is to allow user-generated content: we believe that substantial learning would happen in the course of creating high-quality lessons and glyphs. However, some mechanisms for vetting content would need to be developed as well, since proliferation of low-quality content would severely compromise the player experience.

We have also considered deploying our application to tablet computers, where a finger or stylus could be used to sketch, a more natural action than sketching with a mouse pointer. Deployment to smart phones is another possibility, although the lack of screen real estate and the high degree to which the user's hand would obscure the screen while sketching makes phones a more challenging target. In future work, we also anticipate adding more freeform sketching activities than are currently present, such as unguided sketching of an object's contours, coloring in an object, or using gestural strokes to indicate aspects of large-scale composition or relationships between image elements.

7 Conclusions

This paper described a design for a sketch-based educational game for learning images, tested in the domain of art history. Players had to search through historical paintings for hidden glyphs, and to sketch the outline of each glyph on finding it. Completing a glyph gave access to a small piece of learning content, such as an explanation of something seen in the image or a description of some historical context. The number of hidden glyphs available in the image was shown to the players, encouraging them to continue until all were complete. The design was structured to induce players to spend longer on each image than they would be inclined to in a traditional textbook-style environment; the use of sketching ensured time spent, while lightly occupying the player's mind so as to reduce boredom.

Our user study confirmed that players would spend substantially longer playing the game than studying a textbook, and that the playing time was worthwhile in that it correlated with additional learning. Players' subjective assessment of the game was positive and they viewed it as an appealing alternative to traditional study.

There remains scope for future work. We described additional game elements, including achievements and leaderboards, which would be worthwhile in a long-term deployment. We would like to conduct a longer-term study, probably in the form of providing the game to students to play over the course of an entire semester of an art history course. Lastly, although we tested the design in the domain of art history, we believe that it is suitable for any domain in which images feature heavily, including anatomy and geography. We would like to partner with domain experts in other fields and test our system more broadly.

Acknowledgements

We would like to thank Flickr's Creative Commons for the images in Figure 5. This work received financial support from the GRAND NCE and from Carleton University.

References

ANDRADE, J. 2010. What does doodling do? *Applied Cognitive Psychology 24*, 1, 100–106.

BERTEL, S., BARKOWSKY, T., KONIG, P., SCHULTHEIS, H., AND FREKSA, C. 2006. Sketching mental images and reasoning with sketches: NEVILLE – a computational model of mental & external spatial problem solving. In *Proceedings of the 7th international conference on cognitive modeling*, Edizioni Goliardiche, ICCM '06, 349–350.

EDUWEB, 1997. A. Pintura: Art Detective. http://www.eduweb.com/pintura/. [Accessed 07/04/13].

EGENFELDT-NIELSEN, S. 2011. What Makes a Good Learning Game?: Going beyond edutainment. *eLearn 2011*, 2, 2.

ELECTRONIC ARTS, 2008. Rock band. PlayStation 3.

FISH, J., AND SCRIVENER, S. 1990. Amplifying the mind's eye: Sketching and visual cognition. *Leonardo 23*, 1, 117–126.

GOOGLE, 2011. Google Art Project. http://www.googleartproject.com/. [Accessed 22/04/12].

LOUVRE, 2013. Louvre Museum. http://musee.louvre.fr/. [Accessed 24/05/13].

PAULSON, B., EOFF, B., WOLIN, A., JOHNSTON, J., AND HAMMOND, T. 2008. Sketch-based educational games: "drawing" kids away from traditional interfaces. In *Proceedings of the 7th international conference on Interaction design and children*, ACM, New York, NY, USA, IDC '08, 133–136.

PIPER, A. M., AND HOLLAN, J. D. 2009. Tabletop displays for small group study: affordances of paper and digital materials. In *Proceedings of the 27th international conference on Human factors in computing systems*, ACM, New York, NY, USA, CHI '09, 1227–1236.

SCHELL, J. 2008. *The Art of Game Design: A book of lenses*, vol. 54 of *Morgan Kaufmann*. Morgan Kaufmann.

SOWELL, J. 2009. Learning Cycles in Art History. *College Teaching 39*, 1 (Aug.), 14–19.

TANZ, J. 2011. The curse of Cow Clicker: how a cheeky satire became a videogame hit. *Wired*.