Art Game: An Early Interactive Design from the Office of Charles and Ray Eames

Jim Hoekema, Hoekema Design & Editorial

The Art Game is a little-known but significant early achievement in interactive design. Created in 1977 and 1978, it was one of the earliest examples of interactive multimedia of any kind, and probably the first in humanities education.

Also notable is that the Art Game comes from the legendary office of Charles and Ray Eames. Over the course of their storied career, the husband-and-wife team arguably invented both the interactive exhibit and the multimedia presentation, but the Art Game may be the only Eames “interactive multimedia” in the modern sense of images, sound, and text

linked to individual interaction under computer control. It was also among the last designs completed at the Eames office before the death of Charles in August 1978.

The Art Game is not well known because at first it was top secret. Sponsored by IBM, the project was part of corporate research into possible new market directions, so strict confidentiality was necessary. As it happened, the company did not pursue this direction, and the secrecy melted away, leading to the project’s appearance in *Eames Design*, the compendium of the studio’s work by John and Marilyn Neuhart (Abrams,

1989). To my knowledge, that was the first and only published reference to the project. This will be the first detailed description.

In addition to its claims to be a pioneering interactive design and among the final Eames designs, the Art Game also established a unique approach to developing skills of visual recognition that might still be applied profitably to other areas of image recognition.

THE EAMES OFFICE

The work of Charles and Ray Eames is widely known and well documented. A handy starting point for more information is the Eames website (www.eamesoffice.com), managed by the Eames grandchildren. My own experience there was brief but charmed. As much as I enjoy relating those personal adventures, I will stick here to the project at hand.

My employment at the Eames studio was a happy accident. In August 1977, while on vacation in Los Angeles, I expressed interest in visiting the Eames office to a college friend, Jeannine Oppewall (now a well-known production designer in the movie industry), who had worked there for seven years.

Ray and Charles Eames at an Easter lunch at the office in 1978, a few months before Charles' death.

On my initial visit, after a brief introduction to Charles and Ray, I sat down with the brilliant, Cambridge-educated Jehane Burns. She asked me to “think about doing something about art” using the then new technology of the videodisc. The follow-up to that casual discussion led to an invitation to come in and put the ideas on paper, which we reviewed with Charles. The work was extended in small increments until, aided by Jehane and other designers and filmmakers, we produced a sample sequence of screens. After the client approved the concept, I returned in early 1978 to flesh out the design and create the prototype described below. (Once the project was complete, I stayed on, but Charles died August 21, 1978, and Ray effectively closed the studio at the end of the next year.)

ART GAME: THE PROJECT

Charles Eames was fond of saying that if you analyze a problem in sufficient depth, the design solution will present itself. In this case, the client—an advanced projects group in IBM—had asked Eames to suggest potential consumer applications of a technology that would be commercially released in another year, namely videodisc (aka LaserDisc). This was the large, 12-inch disc that could store and play video—the first incarnation of laser optical technology that later evolved into the audio CD, DVD, and Blu-Ray formats.

Videodiscs had two formats. One could play an hour of running video; the other held less video but could also store up to 54,000 individual frames or images. The industrial player would be able to connect to an external computer, enabling the rapid retrieval of video segments or individual frames under software control. This ability to access thousands of images under computer control was the key capability we hoped to exploit in this project.

Goals of the Art Game. In thinking about doing “something about art,” we reflected that the language of art history and art criticism often frustrates people and seems more barrier than aid to appreciating art. (If you haven’t read an art magazine lately, try one.)

In the Art Game, our goal was to treat looking at art as a learnable skill, separate from the skill of being able to talk or write about the subject. We wanted to break—or at least sidestep—the language barrier in art appreciation,

and to make the essential skill of visual analysis quickly available to any interested adult.

The Art Game doesn’t eliminate “art talk” entirely, but it uses that language when you are most receptive to it, after you’ve been concentrating on the art itself. In the game, you look at the painting first and think about who painted it. You make a choice and then find out if you are right or not. Only at that point do you get some verbal analysis.

As the game goes on, you compare each new painting with others you’ve seen. The longer you play, the more you build on existing visual knowledge to form a holistic mental picture of the style of your selected painter.

Phase 1: Simulated sequence (December 1977). In late 1977, the Eames office produced the initial phase: a series of simulated interactive screens. These simulated screens were intercut with shots of a family interacting with the program as if it already existed in an internal IBM concept film. (The film also included other multimedia concepts, including one on bicycle repair by the MIT group later known as Media Lab.) After reviewing the film, the corporate management committee approved funding for our next stage of work.

Phase 2: Working prototype (June 1978). The Art Game prototype took the design suggested in the sample sequence and produced film segments, images, feedback text, and a computer program so users could actually interact with the design. Created the year before commercial introduction of the videodisc, the prototype was set up using an electronic still store (ESS) device of the type used in broadcast studios, lashed to a refrigerator-size “minicomputer” in an IBM office building in White Plains, New York. A bundle of cables ran from the computer room to a simulated living room, where market research subjects were brought to interact with the Art Game and other prototypes. Subjects were asked questions like “Would you like to have this at home?” and “How much would you expect to pay?” Unfortunately, we were not able to have anyone play the Art Game long enough to see whether their recognition of the painters actually improved!

ART GAME: WALKTHROUGH

Users interacted with the game by means of a simple, custom-built remote control, which contained the key functions for user input.

The game functioned primarily with only three buttons:

- **Next**—Show the next random painting, or the next feedback screen
- **Yes/No**—Is the painting by the painter you chose?

Other functions included:

- **Detail**—Show a close-up view of the current painting
- **Review**—Show the paintings you've correctly identified so far
- **Menu**—Stop studying the current painter and start with another
- **Skip**—Go to the next image without answering (or skip the intro film)
- **Flashback**—Go back to a previous screen (not sure why we didn't just call it "Back")

The following sequence of screen images shows the basic elements of the Art Game prototype, as originally produced at the Eames office.

Title frame. "Art Game" was the working title. For the prototype, we called it "Paintings and Painters" to underscore the nature of the interactions. (In hindsight, the instruction to "push" a button seems a little archaic compared with the gentler "press" that is conventional today.)

Introductory film. A short introductory video explains the concept of the game, calling it "a special way of looking at paintings by trying to distinguish the work of individuals." The narration states: "Your part is to pick out the paintings of one painter. It's a way of training yourself to recognize a painter's style." The film likens the experience to covering up the labels in a museum, "but here ... each time you start down the gallery, the paintings are reshuffled, so you'll probably never see the same sequence twice." The film also explained the buttons on the remote control.

Impressionism film. You are given the option to view a five-minute film that discusses Impressionism and Post-Impressionism, schools of painting important to the game, and the six featured painters.

Select a painter. At this pivotal point, you are asked to select one of the six painters "whose work you would like to know better." Focusing on one painter was a key design decision: The goal was to help users build a mental model of a single painter in a pass, rather than

insisting users try to learn them all at once through a sorting exercise.

Begin random selection. Let's say you select Monet. A screen explains that the system will call up paintings at random from a pool of images. (For more on how that pool was defined, see below.) Your only task is to respond yes or no to the question, "Is it by the painter you selected?"

Random painting display. When you press Next, the system makes a random selection from the pool of images specifically designed for studying Monet. For example, let's say this painting appears:

User input and answer. Let's say you think this is by Monet. You press Yes. The initial feedback screen reminds you of the question (Monet?) and shows your

COVER STORY

input (Yes). After a brief pause, your answer is processed—in this case, you see “Right.” Note that each step was a separate pre-recorded videodisc frame. (In retrospect, the size of the “Right” message seems excessive!)

Yes-Right feedback (Frame 1). After another predetermined pause (not a mechanism I'd recommend today), the screen shifts to the visual-analysis feedback. The size and content of the messages depend on (1) whether this is by the painter you are seeking, and (2) whether or not you correctly recognize it.

Yes-Right feedback (Frame 2). In this case, it was by Monet, and you identified it correctly, so you get a longer feedback. The ellipses in the first frame indicate there is more. The content of the message reinforces your correct response. Pressing Next summons the next frame:

Yes-Right feedback (Frame 3). Pressing Next again brings up the third

and final frame of the analysis feedback (below). In this case, the content includes a quotation by Monet that defines the nature of Impressionism. (In a normal classroom setting, you might invoke this quote early in talking about all the Impressionists, but here, the tidbit could arrive early, late, or in the middle of a user's gameplay.) The red square indicates the end of the feedback: Pressing Next now will summon a new randomly selected painting.

A note on the presentation format. These analysis feedback screens include a small image of the painting, a caption giving the title and date, and some interpretive text. The downside of this format is that users must take their eyes off the image in order to read the text. Also, the low resolution and variable quality of standard television in 1978 meant that text had to be quite large with generous spacing to ensure legibility. That in turn led to a short message like the example above—a mere 55 words—requiring three separate videodisc frames. Of course, we could have made them longer, but we wanted users to spend their time looking at paintings, not reading about them.

During the production of the prototype, we learned of a possibility for a “sound over stills” format, in which audio could be loaded separately and played over a succession of still frames—essentially a voice-over film

but with no “moves” or dissolves in the picture. This promised a great improvement, since we could show the image full-screen while presenting the analysis content in a voice-over. For the feedback shown above, we could cut to a close-up when describing brushwork, for example, and show other Monet Haystack paintings as they are mentioned. We produced 15 feedbacks in this format. This would have been the preferred format, but the news of this possibility came too late to incorporate fully into the prototype.

Regardless of the feedback-presentation format, the content design and information architecture behind the game had some interesting aspects. To that end, let's resume the walkthrough.

Random painting (2). After absorbing the interpretive feedback in the previous interaction, you press Next, and another painting is randomly selected from the pool. Let's say this painting appears:

User input and answer. Let's say you think this painting is also by Monet. You press the Yes button...

Yes-Wrong feedback. In this case, you are wrong. As before, you get the basic

right/wrong feedback, then some visual analysis. In our view, it would not make sense to go into too much depth about a painting you've incorrectly identified. Instead, the analysis offers a quick tip on distinguishing between this work and that of the painter you are studying.

Van Gogh
Fishing Boats at
Saintes-Maries
June 1888

The composition and tonality of the sky and beach recall Monet, but not the strong, flat colors of the boats.

This is by Van Gogh. ■

After a brief analysis, you are back in the game.

Random painting (3). You press Next to summon another randomly selected painting from the pool designed for studying Monet. Let's say this painting appears:

Detail view. Impressionist paintings generally fit the proportions of a TV screen, but the resolution of standard TV was very low. For this reason, as you are considering whether any painting is

by your painter, you can press the Detail button to see a close-up view:

(The prototype offered only a single, pre-selected detail view—a separate videodisc frame. Today's technology would offer more expansive zooming and scrolling abilities.)

Multipurpose feedback (Frame 1).

You are still studying Monet, and you say Yes. That is wrong, because this painting is by Renoir. The feedback begins by acknowledging qualities that Monet and Renoir have in common:

Renoir
Boating Party
at Chatou
1879

Renoir and Monet shared the primary goal of painting momentary effects of light and color in landscape....

Multipurpose feedback (Frame 2).

The feedback then identifies a characteristic that might help you distinguish Renoir from Monet in the future:

Renoir
Boating Party
at Chatou
1879

But unlike Monet, Renoir always liked to include his friends—such as the painter Caillebotte in the foreground here. ■

Actually, this feedback is useful for more than one situation. If you pressed No, correctly realizing this work was not by Monet, the same feedback would

be appropriate as reinforcement. In a third case, if you are studying Renoir, and you mistakenly press No, the same feedback comparing Monet and Renoir helps just as it did for the Monet study.

Same painting, different painter search.

In the prototype, the Renoir painting shown here is also in the pool when you are studying Van Gogh, but in the designers' opinion, this work would be considered pretty far removed from typical work by Van Gogh. For this reason, the feedback is kept short, so the user can get back to thinking about Van Gogh. If you are looking for Van Gogh and respond with Yes at this painting, the game presents the following answer and feedback:

Van Gogh?
Yes
WRONG

Renoir
Boating Party
at Chatou
1879

Renoir was as much a virtuoso of the brush as Van Gogh—but Renoir's strokes end in feathery streaks, Van Gogh's in viscous ridges. ■

The first part of the answer is meant as a kind of compliment—i.e., you got it wrong for this understandable reason. The second part offers a clue to help identify Van Gogh's work in the future.

Note that you may learn about all the painters as you play, but the game guides you to concentrate on a single painter during each pass.

Review function. When considering a new image, you may want to compare it to paintings you've already correctly identified. Some of those may have been

AFTERLIFE

The Art Game existed in fully functional form for only a few weeks in the summer of 1978. It did have a limited afterlife, however, in the subsequent few years.

Impressionists Disc (1980)

IBM redirected its interest in videodisc into a company called DiscoVision Associates, which was licensed to press discs and market an industrial player. DiscoVision produced a demo disc from the assets of the Art Game.

The Impressionists disc included the two video introductions and the 15 sound-over-still segments (feedbacks with audio narration over cuts-only images), but the bulk of the still-frame feedbacks were not included. This was a far cry from the original, but it did provide a means of accessing and displaying some of the material.

Museum Education Consortium (1986–1988)

In 1986 and 1987, a modified description of the Art Game formed the core of two proposals created by a consortium of art-museum education departments, who were interested in applying new technologies in art education. Led by the Museum of Modern Art, the group included the education directors of the Metropolitan Museum (New York), the Art Institute of Chicago, the Philadelphia Museum of Art, Museum of Fine Arts (Boston), the Brooklyn Museum, and the National Gallery (Washington, DC).

My contribution was to adapt the original Art Game design to work on the personal computer. Functions originally triggered by buttons on a remote control could now be enabled through touch controls on the screen, as shown in the sample storyboard (left).

Ray Eames gave her blessing to the project shortly before she died. The consortium was successful in securing funding from the Pew Charitable Trusts, but as it happened, the simple goal of learning painters' styles did not match all the agendas of the multiple institutions. The funded work was split into several parts; in 1988 the project coordinator described the "public access" component as "a non-game approach capitalizing on search and sort capabilities." The Art Game was used for some user testing but not further developed.

Art Game HyperCard Stack (1988–1990)

Around this time—August 1987—Apple released HyperCard, an interactive toolkit that also had the ability to control a videodisc player! Thanks to its simple coding language, even I, a non-programmer, was able to re-create the logic of the Art Game in HyperCard, complete with random selection, pool maintenance, and so on. Most of the content appeared on the Macintosh screen, but the program triggered certain still frames or video segments to play on an attached videodisc player running the Impressionist videodisc. One version was designed to access only paintings that had the sound-over-still [voice-over audio] segments.

While the game's logic did not change, the new platform again provoked design shifts. A mouse now delivered user input, and the feedback fit on the same screen as the black-and-white image. (A color image from the videodisc would be on the TV.) This version also added features like a running score and counters for the selection and review pools.

The museum education consortium tested this version of the Art Game with visitors at the Museum of Modern Art in September and October 1988. The report cited a generally positive response, with several commenting that it was "a fun way to learn about art," but the study was based on a sampling of only seven self-selected museum visitors.

What's Next?

The Art Game was an exploratory design at the Eames office almost 40 years ago, but I've never seen anything quite like it in the years since. The concept survived changes of delivery format, and I can't help thinking it would survive well—with changes—on a mobile platform. It seemed to us at the Eames office that the game's strategies for content, logic, and feedback might work, not just for other periods of art, but also for other types of visual-recognition tasks, such as medical imaging, aerial photography, microphotography—almost any kind of imagery that comes with enough samples and in a manageable number of classes with distinct characteristics.

In the meantime, I thought you should know about this interesting project originating with one of the world's great design practices back in the day!

lucky guesses, but now you may want to look again. You can do that by pressing the Review button:

Images that the user already correctly identified appear. You can walk through them by pressing Next, perhaps calling up a Detail for a closer look. While your initial recognition may have stemmed from some obvious clue, the hope is that when looking again, you may notice other subtleties that will add to your mental image of the painter you are learning to recognize.

Exit strategies. There is no score in the Art Game. In theory, the game does not end until the user can recognize all the available images of each painter. That could happen easily enough in the prototype, which had only 15 images per artist. But it seemed harsh to force users to continue to the bitter end, so the game presents a gentle reward message at an earlier point, as follows:

(We did not hardwire what a “good number” should be—that tipping point could change based on the scope and difficulty of the paintings in the game.)

Still, some users may want to keep going to that bitter end, for closure, or just because they can. If that were

you, you would eventually receive a message politely explaining that you have reached the limit of the system’s knowledge:

As the final message suggests, you could start over on the same painter, and it would be a different experience if only because the images would randomly appear in a different order. If you now elect to study a different painter—say, Cézanne—some of the images by Cézanne and others will be familiar if they were in your previous Monet pool, but there will be new works by Cézanne, and new wrong images as well that you would not have seen in your initial painter study. Thus, despite the random element, your knowledge should grow cumulatively—the last painter you study should come easier than the first, regardless of the order you selected.

As shown in the walkthrough, the Art Game offered a simple pattern of interaction, but it was designed to enable some fairly sophisticated learning mechanisms, thanks to a design approach that integrated expert knowledge with a clever approach to logical structure.

LOGICAL STRUCTURE

The following sections describe the logical structure and content strategy that underlie the Art Game experience.

Weighted random selection. When you are ready for a new painting, pressing Next selects one image at random from a pool of images, but that pool is customized for the painter you are studying. If selection were truly random, and there were six painters in the game, you’d get a Monet (say) on average once in every six images; but because you are presumably learning this material, you’d make some

mistakes. If you spotted only every other Monet correctly, you would enjoy the satisfaction of a right answer only once for every 12 images! Spot one in three, and you’d be right only once in 18 attempts! That would be discouraging.

Plus, there is another problem. From the designers’ point of view, “open” random selection would mean that we would have to create a specially crafted feedback for each possible outcome. That could mean six or even 12 different feedbacks to cover the Yes and No states for all six painters. This was daunting—we did not want to write that many different descriptions of the same painting.

For those reasons, the random selection is taken from a weighted or tailored pool consisting of all the paintings by the target artist, plus about twice that number from the remaining five painters. Thus, when you select Renoir, for example, the pool from which your random images will be drawn is defined as all the Renoirs we have, plus roughly twice that many from the other five painters. Now, for a given random selection, you will have a roughly 33 percent chance of seeing a Renoir, which should make game play more satisfying.

From a data perspective, this means that each painting is a candidate for selection when that painter is under study, but it will also have “flags” or logical attributes indicating when it can be selected as a wrong answer under other painter searches. For example, a painting by Monet will turn up in the Monet pool but might also appear in the pools for Renoir and Van Gogh. Typically, each painting is flagged for two other pools, although the prototype had a few exceptions.

Updating the pools. Another bit of logic behind the scenes has to do with how the pools are updated. When you input a correct answer, that painting is removed from the selection pool; if

ACM Transactions on Reconfigurable Technology and Systems

This quarterly publication is a peer-reviewed and archival journal that covers reconfigurable technology, systems, and applications on reconfigurable computers. Topics include all levels of reconfigurable system abstractions and all aspects of reconfigurable technology including platforms, programming environments and application successes.

www.acm.org/trets
www.acm.org/subscribe

Figure 1. Art Game logic of random selection and pool maintenance.

A04 Monet: La Grenouillère. 1869. 29 x 39"
Metropolitan Museum of Art, New York. (MMA slide)

Figure 2. Art Game example feedback pathways.

you answer incorrectly, the image goes back into the pool. For example, if you are studying Gauguin, a painting by Gauguin appears, and you correctly select Yes, then that painting is removed from the selection pool. The reasoning is that if you have correctly attributed the image, there is no need to ask you again. Correct attributions also go into the Review set, where you can access them again as needed.

Similarly, if you are studying Gauguin, a Renoir appears, and you correctly say that it is not a Gauguin, that image is also removed from the selection pool. If you have demonstrated the knowledge that this is not Gauguin, there is no need to repeat the question.

On the other hand, if your response to either image was wrong (not recognizing the Gauguin as such, or mistaking the Renoir for a Gauguin), then that image goes back into the selection pool, where random selection will sooner or later retrieve it again for your consideration. If the learning process is working as designed, you'll get it right the next time the image appears!

The logic of random selection and pool maintenance is shown in Figure 1.

It should be noted that the allocation of paintings to pools was not calculated using any sort of programmed algorithm. This was a purely handcrafted process, the result of the visual analysis of the paintings by art-historian author-designers.

CONTENT DESIGN

We liked to think the Art Game had a general-purpose engine for learning visual content of many different kinds, but we must confess that the topic we selected for the prototype was particularly well suited to the initial proof-of-concept prototype.

Why Impressionism? The Impressionists painted ordinary everyday scenes, easily understood by anyone. Among the many periods of art history, this is one where you don't need to identify the names of saints, gods, and goddesses, or learn obscure stories in order to appreciate the work. Largely a school of landscape artists, Impressionism tends to fit well in the

rectangle of a television screen. Charles Eames noted ironically that Cézanne especially seemed to hew closely to academy ratio, the old 3x4 ratio of all movies before the wide formats came along. Also, the Impressionists and Post-Impressionists primed their canvases in white and painted with bright colors, so their paintings come through on a TV better than those of earlier periods.

Selecting the painters. The set of painters included in the game essentially determined its level of difficulty. The six painters featured—Monet, Renoir, Cézanne, Van Gogh, Gauguin, and Seurat—had distinctive styles, yet they shared goals and painted many similar scenes. Anyone who has taught or taken an art history course knows that the best way to study the style of one work is to compare it to another rendition of the same subject by another artist. With that principle in mind, we selected mostly landscapes, plus a few still life and portrait images, in order to give each artist an even footing in the random selection that drives the game's structure.

We considered this collection of three Impressionists and three Post-Impressionists as a fairly easy challenge for people with little exposure to art appreciation. A different set of painters would result in a different level of difficulty.

Selecting the paintings. The feedback examples above showed that you often learn the most from an “intelligent wrong guess,” where the feedback offers insights both on the current image and on the painter you are studying. For us as designers, it was tempting to include many “ringers”—works by one painter that look like another's. Examples include an exceptional unpeopled seascape by Renoir that looks like a Monet; a later Monet landscape as fevered as a Gauguin; any of the early landscapes by Van Gogh or Cézanne that look as Impressionist as the work of Monet or Renoir; and so on. But if the game contained too many trick paintings and not enough typical works, users would never develop a reliable model of each painter's characteristic style. The prototype therefore contained a majority of typical works along with a smattering of those that make for learning moments or interesting wrong guesses. Domain

knowledge was essential to data definition.

Writing the feedbacks. A typical painting in the game appears in the pools for three painters, allowing for six different interaction paths. But it was difficult to devise even that many distinct targeted messages, so most paintings have only three feedback messages, which are deployed to serve the six different pathways. An example of these pathways is shown in Figure 2.

This diagram shows the target pattern, where one long answer reinforces a correct attribution; a short answer addresses a “dead wrong” answer; and all other paths receive the “intelligent wrong guess” medium-size feedback.

Writing these feedbacks required discipline to ensure they would make sense and sound appropriate in any of the answer pathways and at any point in gameplay. The feedbacks are all self-contained, with no references like “as we saw earlier.” Due to random selection, a given message might be the first or the last feedback a user encounters.

CONCLUSION

The Art Game offered a very simple user experience; the bulk of the gameplay consists of answering Yes or No to each painting as it's presented. But attribution is a non-trivial question—expert connoisseurs are concerned with exactly that when authenticating a discovery. Like them, Art Game players bring their knowledge and instincts (however recently gained) to bear on that one question.

As with all humanities, art appreciation is not a cut-and-dry subject. We didn't try to turn connoisseurship into a science, but rather focused on one binary factor—authorship—to help users improve their skills in visual thinking and recognition.

Jim Hoekema has been an interactive multimedia and user experience designer since the project described here. He has worked independently and for firms including Philips Electronics and Accenture/Fjord. Hoekema Design & Editorial is his independent consulting firm in Newburgh, NY.
→ jim@hoekema.com