Simulating Visual Systems using NPR Techniques: Methodology, Framework, and Case Studies

by

Jeremy Long
B.Sc., University of Saskatchewan, 2005
M.Sc., University of Saskatchewan, 2007

A Dissertation Submitted in Partial Fulfillment of the
Requirements for the Degree of

DOCTOR OF PHILOSOPHY

in the Department of Computer Science

University of Victoria

Simulating Visual Systems using NPR Techniques: Methodology, Framework, and Case Studies

by

Jeremy Long
B.Sc., University of Saskatchewan, 2005
M.Sc., University of Saskatchewan, 2007

Supervisory Committee

Dr. A. Gooch, Supervisor
(Department of Computer Science)

Dr. M. Tory, Departmental Member
(Department of Computer Science)

Dr. B. Wyvill, Departmental Member
(Department of Computer Science)

Dr. T. Pelton, Outside Member
(Department of Curriculum and Instruction)

Supervisory Committee

Dr. A. Gooch, Supervisor
(Department of Computer Science)

Dr. M. Tory, Departmental Member
(Department of Computer Science)

Dr. B. Wyvill, Departmental Member
(Department of Computer Science)

Dr. T. Pelton, Outside Member
(Department of Curriculum and Instruction)

Abstract

In this dissertation I examine how research in non-photorealistic rendering, human perception, and game-based learning can be combined to produce illustrative simulations of different visual systems that effectively convey information about vision to unprimed observers. The Visual Differences Simulation (VDS) methodology and framework that I propose is capable of producing simulations of animal visual systems based on how they relate to human vision, and can represent differences in color vision, hyperspectral sensitivity, visual acuity, light sensitivity, field of view, motion sensitivity, and eye construction. The simulations produced by the VDS framework run in real time, allowing users to explore computer-generated environments from 'behind the eyes' of an animal in an interactive and immersive manner. I also examine how cognitive principles and game-based learning can be leveraged to demonstrate and enhance the educational impact of the simulations produced by the VDS framework. Two case studies are presented, where simulations of the cat and the bee visual systems are used as the basis to design educational games, and are evaluated to show that embedding the simulations in educational games is an effective and engaging way to convey information about vision to unprimed observers.

Contents

List of Tables

List of Figures

Acknowledgements

First of all, I would like to thank my supervisor and supervisory committee for the comments, suggestions, and guidance they have given me throughout this process. I would also like to acknowledge the department graduate secretary, who helped me stay on track throughout the administrative steps in this process.

I would also like to acknowledge my colleagues in the graphics labs, particularly those that worked with me on some facets of these projects. In particular, I would like to thank Anthony Estey, Sven Olsen, and David Bartle for their contributions to the *Catalyst* project. I would also like to acknowledge all the participants in my experiments who kindly donated their time on the altar of research.

Last but not least, I would like to thank my family and friends for keeping me walking this road until its conclusion. In particular, I would like to thank Donna for all the patience and help she offered throughout this journey. I could not have traveled half so far without your guidance and support.

Chapter 1

Introduction

In this dissertation I propose the Visual Differences Simulation (VDS) methodology and framework that can be used to produce real time illustrative simulations of animal vision, and examines how cognitive principles can be employed to design educational games that demonstrate and enhance the impact of the simulations. The VDS framework allows users to explore a virtual environment from 'behind the eyes' of another species in a more immersive and interactive manner than was previously possible. The educational impact of this approach is demonstrated by two case studies, where the simulations of the cat and bee visual systems are incorporated into educational games. User studies show these games and the simulations they contain are an effective and engaging way to convey information about animal vision to unprimed observers.

Understanding how an animal species sees the world is an important step in understanding that species' behavior. The bumblebee can serve as an example of how visual characteristics can influence behavior. Despite having relatively few photoreceptor units in their eyes, bumblebees are capable of recognizing and identifying different types of flowers, and pollinating accordingly. They cannot see shapes at a distance as well as humans, so they rely instead on other senses and visual characteristics such as their color vision system [16]. Many flowers have patterns of pigment on their petals that reflect ultraviolet light that is invisible to humans, but can be detected by bees. Extended spectral vision is just one example of how a species' visual system has evolved to help it survive and thrive in the wild, and how understanding bee vision enables us to better understand the natural world.

Neuroscientists and biologists have been studying animal visual systems for decades. However, they tend to represent their results in the form of diagrams, graphs, and spectral sensitivity curves. These are meaningful artifacts to those in the discipline, but they are not likely to give the average person a complete picture of how an animal species sees the world. At the same time, there is concern among educators that younger generations are growing up in greater ignorance of the natural world, and that this could have long-lasting consequences on conservation and education [30]. Educators and interpretive centers, such as zoos and museums, can take advantage of new media and technology in order to better engage young audiences, and give them a greater appreciation for animals and the research that is being done to better understand them. After all, protecting animal species for the future can only succeed if future generations are interested enough to protect them.

I contend that simulating visual systems is a domain that benefits from a visual representation. Research in non-photorealistic rendering (NPR) suggests that illustrative images can be an effective, and often compact, way of communicating information [15, 26]. This notion can be traced back to master artists, who used artistic techniques in attempts to portray more than just the lines on the canvas, but the meaning, passions, and feelings their works of art represent. Artists and illustrators have also employed abstraction and simplification to produce sketches, maps, and scientific illustrations that reduce the content of images to the bare essentials most useful to the task at hand, increasing the ratio of important information perceived by observers [1, 26, 27, 66, 67, 78]. These techniques are motivated by the idea that non-photorealistic images can sometimes better convey pertinent information than photorealistic images, text, or audio.

Creating a visual representation of animal vision is complicated by the fact that some of the information that should be included is visually ambiguous from a human perspective, such as hyperspectral sensitivity. This information needs to be embedded into the visual display in a meaningful manner, and I accomplish this using a *difference-based* simulation methodology inspired by research in NPR and human perception that represents visual systems based on how they relate to human vision. The VDS framework employs the *difference-based methodology* to represent a wide array of visual characteristics including color vision, hyperspectral sensitivity, light sensitivity, visual acuity, motion sensitivity, field of view, and eye construction. The framework takes advantage of modern graphics hardware and algorithms to produce simulations that run in real time. The level of immersion and interaction offered by these simulations goes beyond the 2D image filters that have previously been used to visualize aspects of animal vision [36, 83].

I also examine how game-based learning can be leveraged to demonstrate and enhance the educational impact of the simulations produced by the VDS framework.

Experiential learning is of particular relevance to zoos, museums, and other interpretive settings, where there is increasing interest in using *edutainment* applications to supplement existing exhibits and engage young people [42, 85]. Games based on simulations could also be used in schools to reinforce the science curriculum—to help students learn more about the human visual system by contrasting it to the way that other species see the world.

Chapters 11 and 12 describe two case studies where the simulations produced to represent cat and bee vision are integrated into educational games that were designed using cognitive principles. The educational games were evaluated through user studies that show their efficacy in conveying information about the visual differences between humans and the respective animal.

In summary, this dissertation contains three major contributions:

- Establishes the importance of a difference-based simulation methodology for representing characteristics of an animal visual system relative to human vision. The VDS methodology draws inspiration from work in non-photorealistic rendering (NPR) and human perception, and could be employed in other contexts where multiple channels of information need to be combined on a single interactive display.

- Describes a framework that builds upon the difference-based methodology to generate simulations that illustrate a wide variety of visual characteristics, including color vision, hyperspectral sensitivity, light sensitivity, visual acuity, motion sensitivity, field of view, and eye construction. The simulations produced by the VDS framework give users the opportunity to explore a virtual world from ‘behind the eyes of an animal’ in a more immersive and interactive manner than has previously been achieved.

- Documents and evaluates two educational games that were created to incorporate and enhance the cat and bee simulations produced by the VDS framework. Both games were designed using cognitive principles, and have proven effective in educating and engaging players with regard to animal vision. The process used to design and evaluate these games is based on and extends existing research in the field of game-based learning, and can help involve players in more immersive learning experiences. These games have potential to be integrated into educational and interpretive settings such as zoos, museums, or even as supplements to the science curriculum in schools.

Chapter 2

Terminology and Background

The VDS framework I propose in this dissertation is based on research in three domains: 1) biology and neuroscience research devoted to human and animal vision, 2) rendering techniques from NPR and computer graphics, and 3) pedagogical principles from game-based learning research. I begin by detailing the characteristics of human and animal vision that are suitable to be simulated in a visual manner, and how they are informed by physiological factors. Next, I discuss attempts that have been made to simulate the different visual characteristics from the artistic and computer graphics communities. Finally, I conclude with a survey of research in game-based learning and the pedagogical principles that make it effective for conveying information through engaging interactive experiences.

2.1 Visual Systems

Visual systems include a complex set of biological mechanisms that process light patterns into information useful to an organism [43]. The visual system builds *visual perceptions*—what an organism actually sees—based on the signals generated by its sensors and receptors. Several biological and neurological processes contribute to this task, including light reception, combining information from multiple projections, and the identification and categorization of visual objects, to name only a few. In fact, some of the factors that contribute towards building a visual perception are not yet entirely understood [43, 84].

The simulations described in this dissertation are not intended to convey knowledge of the physiological elements at work in visual systems. Instead, the goal is to illustrate some aspects of their collective output—the visual perceptions that are produced. Consequently, the emphasis of this section is on identifying visual characteristics that can vary between visual systems, and on examining how anatomical factors identified by researchers can contribute towards these differences. These anatomical factors are used as input to the VDS framework in order to generate simulations that illustrate differences in color vision, light sensitivity, field of view, and a variety of other characteristics.

Photons of light act as input to the visual system by interacting with photoreceptors—cells that contain light-absorbing chemicals, and generate a neural signal when triggered [43]. Rods and cones are the two types of photoreceptors present in human eyes. Rods are more sensitive to light, and are thus dominant under low light conditions, while cones are less sensitive to light and are dominant under brighter conditions [84]. The exact nature of this anatomical framework and the distribution of rods and cones have a significant impact on several visual characteristics that contribute towards the production of a visual perception, including color vision, light sensitivity, and visual acuity [43].

Color sensitivity is one of the most studied visual characteristics, and is mediated under bright light conditions by the number and types of cone photoreceptors present in the visual system. The human visual system typically includes three classes of cone photoreceptors, each containing pigment sensitive to different wavelengths of light. The colors that we perceive are determined by comparing the amounts that the three cone classes are stimulated. Stimulation of the short wavelength cones (S-cones) in isolation gives the appearance of blue. Stimulation of the S-cones and the long wavelength cones (L-cones) gives the perception of purple. Stimulation of the medium wavelength cones (M-cones) with minor stimulation of the L-cones gives green, while stimulation of the L-cones in isolation gives the appearance of red [84].

Color vision deficiencies can occur when the pigment in one or more cone classes is not sufficiently distinct from the others, compromising a person's ability to discriminate between certain colors. Several different strains of color blindness have been identified, depending on which cone class is compromised. Deficiencies in the L-, M-, and S-Cones are referred to as protanopia, deuteranopia, and tritanopia, respectively [62]. The former two conditions are most common, and result in difficulties in distinguishing between red and green colors. The colors seen by those suffering from protanopia or deuteranopia are often visualized as combinations of blue and yellow [65].

The three cones classes that humans typically possess are said to make us a trichromatic species. This is relatively rare among mammals [65]. Many mammals are dichromats, and only possess two distinct cone classes. Dogs, horses and cats are thought to fall into this category, possessing two cone classes with relatively similar sensitivities to the S- and M-cones in the human visual system, likely giving them color perception akin to humans suffering from protanopia or deuteranopia [55, 68]. Species such as bees, pit vipers, and birds have cones with sensitivity outside the human spectrum, allowing them to perceive ultraviolet (UV) or infrared (IR) wavelengths of light [70, 73].

The density and distribution of rod and cone photoreceptors across the eye is also an important factor for light sensitivity and visual acuity [43]. Visual acuity refers to the sharpness of the perceptions produced by the visual system, while light sensitivity refers to the amount of light necessary to stimulate the visual system. A greater concentration of rods will grant the visual system a lower minimum light detection threshold, allowing it to function in darker conditions, but produces a more blurry perception. A higher concentration of cones makes the visual system less sensitive to light, but can increase color sensitivity and visual acuity. These two characteristics are related in other manners as well. Some species, such as cats, have a reflective layer at the back of their eyes called a *tapetum* that bounces unabsorbed light back into the photoreceptors, increasing light sensitivity at the expense of scattering the light and thus decreasing visual acuity [6, 43].

The placement and construction of a species' eyes can also have a significant impact on visual acuity, in addition to other characteristics such as field of view, depth perception, and motion sensitivity. Humans have frontally-placed eyes, each with a single lens that deforms in order to change the focal distance of the eye. This configuration allows for considerable overlap between the eyes, and our depth perception benefits from this region of binocular vision [43]. Some species have laterally-placed eyes that offer a much wider field of view, but at the expense of reduced overlap between the eyes, compromising depth perception. For example, horses are thought to have nearly monocular vision [3]. Insects such as flies and bees have compound eyes, with thousands of overlapping receptors and lenses, each aimed in a slightly different direction [17, 82]. This can offer a very wide field of view, but limits the total number of photoreceptors, reducing visual acuity.

Biologists and neuroscientists continue to research the visual systems of various animal species. Table 2.1 shows a selection of the data that has been gathered. Some of the data in Table 2.1 is still speculative, but is included here to illustrate the wide diversity of visual characteristics that exist in the natural world. The table demonstrates that the characteristics discussed above are far from the only ones that are important in informing a species' *visual perception*, and that each species has a visual system that is suited to their situation. This lends weight to the idea that vision is linked to evolution and behavior, and that each species has developed a visual system that suits their needs.

The nature of a species' visual system can be linked with its behavior, and it has been suggested that learning more about one can help us understand the other [16]. Evolutionary pressures contribute to the formation of a species' visual system, as mutations that allow a species to survive and thrive are more likely to be propagated to future generations [43]. The formation of the human visual system follows this pattern. We have evolved as a diurnal species, more active when light is plentiful. Our visual system is consistent with this lifestyle, and is more attuned to visual acuity than to light sensitivity [43]. The L-Cones in our visual system are thought to have developed because they offered an evolutionary advantage in foraging for edible food [65]. Similarly, the dense concentration of cone cells that makes up our foveal region is thought to have evolved to offer enhanced acuity in our central gaze [43]. This allows us to better detect high frequency details in areas where our vision is focused, which also served to compliment our foraging capabilities.

2.2 Simulating Visual Systems

Neuroscientists and biologists have researched visual system for decades, both to attain a greater understanding of the physiological and neurological processes at work within them, and to gain more insight into animal behavior and evolution [16, 55, 68]. The results of their work are typically depicted as figures, graphs, and diagrams, which are meaningful to those within the discipline, but do not give the 'full picture' to the average person. There have nonetheless been a few notable attempts to represent this information in a more visual manner.

Williams et al. [83] looked at visually simulating insect vision as part of a photographic filter system. Their method used a physical mechanical lens to simulate the compound vision of bees and other insects, and then applied post-process filter effects to represent the color vision and hyperspectral sensitivity of bee vision [17, 82]. Figure 2.1 shows some of their results. However, their process did not run at interactive rates, and functioned only to produce two-dimensional output. These factors limit the observer's ability to explore the visual system being simulated.

Species	Color Sensitivity	Acuity, Light Sensitivity	Field of View	Other Characteristics
Humans	• Trichromatic, cones: 430, 539, and 572 nm. • Red cone offers advantage in foraging.	• Favors acuity over light sensitivity. • Dense foveal region increases acuity.	• Favors depth perception over field of view. • Frontal eyes.	• Trichromatic vision is rare among mammals. • Dense foveal region is effective for foraging.
Cats	• Dichromatic, cones: 450 and 550 nm. • Color sensitivity like humans who are red-green color blind.	• Favors light sensitivity over acuity. • Reflective layer (*tapetum lucidum*) bounces light back to receptors, giving a second chance to detect it.	• Favors field of view compared to humans. • Frontal eyes. • Smaller overlap between eyes than humans.	• Concentrated band of photoreceptors known as the ‘visual streak’ increases sensitivity to motion.
Bees	• Trichromatic, cones: 360, 450, and 520 nm. • Can detect ultraviolet wavelengths of light.	• Favors field of view over acuity and light sensitivity. • Compound eyes include thousands of partially overlapping photoreceptor units, with limited acuity and light sensitivity.	• Favors field of view compared to humans. • Frontal eyes. • Compound eyes allow a wider field of view.	• Extra light receptors on top of heads called *ocelli* that offer additional light sensitivity. • Many flowers reflect ultraviolet light that bees can detect.
Pit Vipers	• Dichromatic, cones: 430 and 550 nm. • Some snakes are sensitive to infrared light, which is combined into a single visual perception.	• Favors light sensitivity over visual acuity. • Allows them to function in dark conditions, in conjunction with infrared sensitivity.	• Favors field of view over depth perception. • Lateral eye placement. • Pit vipers have around 250 degree field of view, with only 35 degrees of binocular overlap.	• Vision is weak, but is configured to promote motion sensitivity.
Horses	• Dichromatic, cones: 430 and 540 nm. • Little sensitivity to light we perceive as green and red.	• Favors light sensitivity over visual acuity. • *Tapetum lucium* offers a second chance to detect light. • Concentrated band of cone cells that boosts acuity and motion sensitivity.	• Favors field of view over depth perception. • Laterally placed eyes. • Very wide field of view at the expense of depth perception.	• Lateral placement of eyes leads to a blind spot in front of face, can be mitigated by tilting head.
Avians	• Some are tetrachromats, with four cone peaks. • Can potentially distinguish colors that look identical to humans.	• Generally optimized for visual acuity. • Nocturnal birds sacrifice acuity for enhanced light sensitivity.	• Predators have frontal eyes promoting depth perception. • Prey species tend to have lateral eyes with wide fields of view.	• Predators prioritize motion tracking to help catch their quarry. • Some have oil drops in their eyes to reduce haze for distance vision.

Table 2.1: A selection of animal visual characteristics gathered from a variety of sources [51, 55, 65, 68, 70, 73, 84]. Note that this data varies between species, and even among individuals within a species. Furthermore, some of the characteristics listed in this figure are speculative, and have not been conclusively determined. They are included here to illustrate the wide diversity that exists in the natural world.

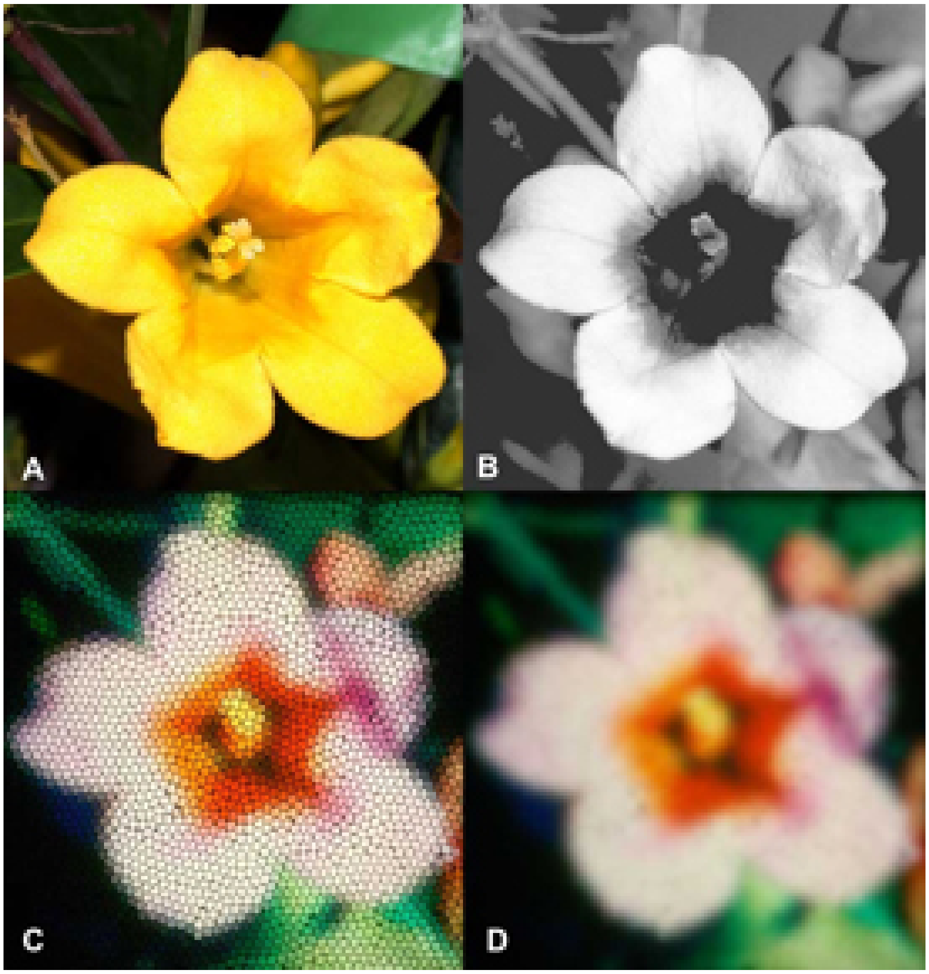

Figure 2.1: The bee visual simulation produced by Williams et al. [83] included four effects. An input image [a] was combined with hyperspectral information [b] using a false coloring transformation. The result was viewed through a mechanical lens structure to simulate bee compound vision [c]. Finally, a blur effect was applied to produce the final image [d]. This simulation required a physical contraption for image acquisition, and did not provide an interactive representation of bee vision. Image used with permission.

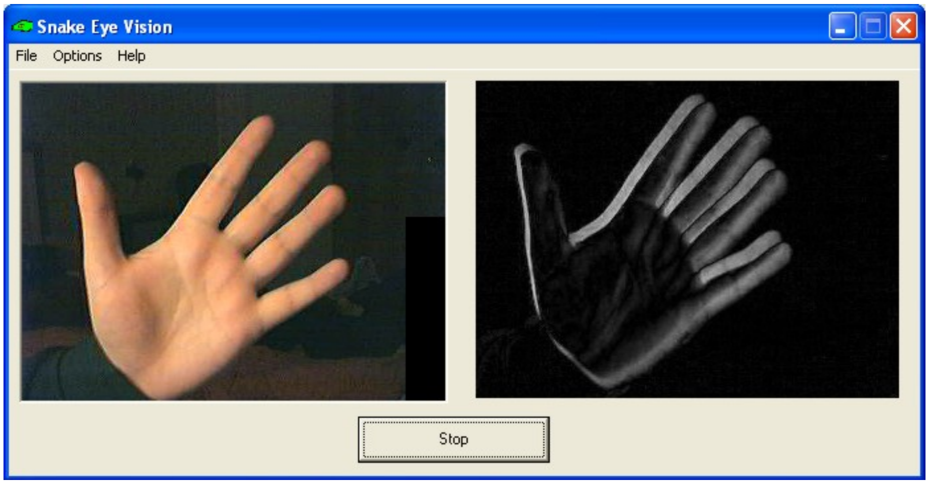

Figure 2.2: Snake Eye Vision [11] is free software produced by Cognaxon that was developed to demonstrate how snakes see the world. The software transformed a webcam feed to show aspects of snake vision, but it attempted to show a small set of visual characteristics, such as motion tracking. The simulations proposed here are more expressive, capable of representing many more characteristics of snake vision. Image used with permission.

Giger [23] built an online image processing application that simulated some aspects of bee vision in an interactive manner—most notably the effect of the bee's compound eyes. However, his approach did not attempt to represent several other important visual characteristics, such as bee color vision, and hyperspectral sensitivity.

Snake Eye Vision [11] is free software that was developed by Cognaxon to demonstrate how snakes see the world, as shown in Figure 2.2. The software transformed a webcam feed to show aspects of snake vision, but it only attempted to simulate a small set of visual characteristics, such as motion tracking. The illustrative simulations proposed here are more expressive, capable of representing many more characteristics of snake vision.

The ZooMorph project [36] began in 2009, and is aimed at developing image filters and plugins to simulate different types of animal vision. Figure 2.3 shows an early rendering of what the project is aiming to produce. Once again, this approach is based on image filters, allowing only limited user interaction.

There has been some research done in the computer graphics community that can be applied to simulate different visual characteristics such as color vision, compound vision, motion sensitivity, visual acuity, and hyperspectral sensitivity, although it was not always intended for this particular purpose. Previous work has generally examined each characteristic in isolation, and has not attempted to show them in context with each other.

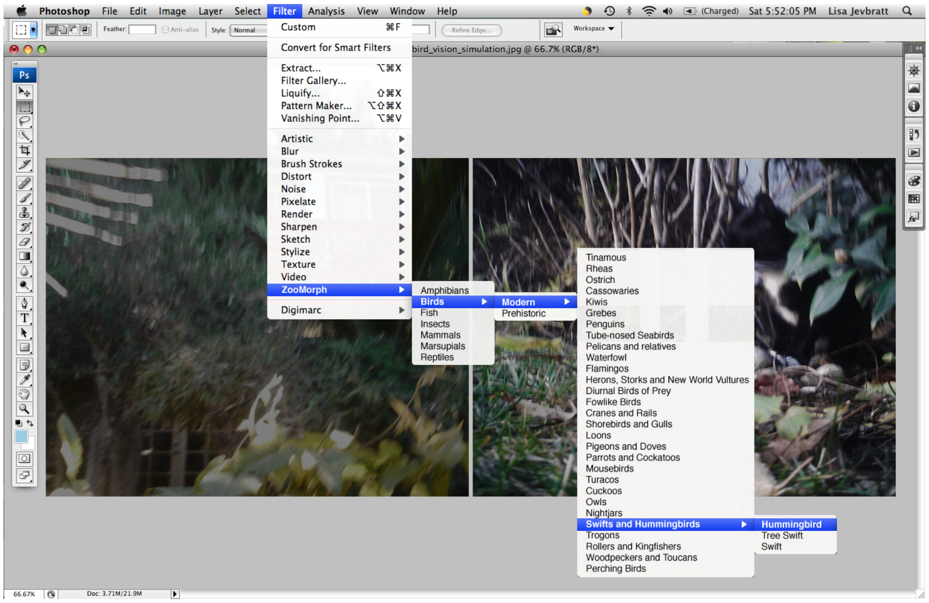

Figure 2.3: The ZooMorph project [36] intends to develop image filters to show how different animals see the world. They seek to simulate animal vision using off-line, two-dimensional image processing techniques. Image used with permission.

Color vision is one of the aspects of the human visual system that has been examined most extensively by computer scientists. Researchers have sought to create color transformations that can simulate various color vision abnormalities [7, 52, 63]. Filters have been developed for this purpose, and can be used both to teach people about different strains of color vision, and also to serve as a basis to create imagery that is more inclusive of color vision deficiencies.

A common approach to this problem involves building a color space based on cone responses, and then collapsing one or more of the dimensions of the color space to a constant value based on which cone class is absent for a particular visual disability [7, 52, 80]. Rasche et al. [62] proposed an alternate method, aimed specifically at preserving contrast and detail when colorizing grayscale images for observers with anomalous color vision. Ichikawa [35] proposed an approach where genetic algorithms are used to help recolor images so as to preserve detail and minimize distance between input colors and their corresponding remapped color. Ma et al. [47] used a self-organizing map (SOM) algorithm to create a nonlinear color map that maintains the neighboring relations between colors. Machado et al. [48] proposed a physiologically-based model derived from electrophysiological data and evaluated their approach with user studies.

Tools have been developed based on these algorithms that show designers how color palettes would look to those with color impairments, and can help the designers select palettes that are more inclusive of color vision deficiencies [72, 80]. Several of the color picking tools that have emerged suggest the idea of emphasizing differences as a way of illustrating visual information. They allow users to toggle freely between the color space for normal human vision and the color space for those with visual abnormalities, helping users build a mental mapping between the two. The approach presented in this dissertation for simulating color vision is also based on illustrating differences by contrasting them with a known human perspective. However, the VDS framework provides a more immersive and interactive experience, simulating differences in *visual perception* in real time, allowing the user to explore the color sensitivity being simulated in context with other visual characteristics, such as acuity, field of view, and eye construction.

Computer scientists and perception researchers have also attempted to convey information about the human visual system using visual media such as images and video. To et al. [79] considered how objects viewed by peripheral vision differ from ones observed in the foveal region, and Raj and Rosenholtz [60] presented a method for visualizing how stimuli appear during peripheral viewing. Barsky [4] used wavefront data from human subjects to produce vision-realistic images consistent with a particular individual's visual system, while Deering [13] produced a photon accurate model of the human eye. My goal is similar to some extent, except rather than attempting to represent aspects of the human visual system in isolation, the framework proposed here is intended to represent the visual perceptions produced by other visual systems in an illustrative manner.

Several methods have been proposed for achieving non-linear perspective, representing the idea of combining the views from multiple cameras with different characteristics into a single view [12, 76]. These approaches could be used as a means of representing compound eye construction, which can include multiple overlapping perspectives of the same scene. However, these approaches do not run in real time when tasked with handling the hundreds or thousands of separate views typically found in compound visual systems.

Researchers have also considered how to enhance or emphasize motion in computer-generated imagery and environments, and these techniques could potentially be deployed to represent the different motion sensitivities that exist in visual systems. Motion blur has been employed in 3D rendering to make motion appear less staggered and more smooth, as it takes into account the integration that cameras capture over the course of their exposure time when recording movement [59]. However, motion blur can make it more difficult to ascertain the current position of moving objects, as they become more blurred.

Instead of attempting to make virtual movement appear more realistic, some researchers have sought to emphasize motion using techniques inspired by animation [32, 38, 50]. Masuch et al. [50] built a system for embedding speed lines in images, where the lines followed moving objects and indicated the direction of motion. Kawagishi et al. [38] looked at 'ghosting' techniques, where moving objects left behind a trail of repeated images that fade away over a limited life span. Haller et al. [32] implemented a 3D system that incorporated speed lines, repeated images, and squash-and-stretch as means of depicting motion. While all three techniques seem to be successful in enhancing motion, the speed lines and squash-and-stretch approaches sacrifice more realism and are more appropriate in cartoon settings. A repeated image approach is more conducive to representing enhanced motion sensitivity without sacrificing realism.

2.3 Game-Based Learning

The purpose of the VDS framework is to produce simulations that can illustrate differences between visual systems to unprimed observers. As such, the VDS framework is only effective insofar as it is able to convey the correct information to the viewer. Advances in graphics hardware and techniques have made it possible to visualize and interactively explore data within a 3D computer-generated environment in real time, and this paradigm shows potential for involving the user in a more immersive and interactive learning experience.

Recent research has suggested that using immersive computer-generated environments in an educational context can prove fruitful. More and Burrow [54] adopted virtual environments from video games for use in architectural design studios, and found that this approach offered advantages in terms of interactivity and immediacy. Johns and Shaw [37] explored the idea of using real time immersive environments for collaborating on the design cycle, from conceptualization to prototyping.

The consensus that follows from this research is that using a computer-generated environment is not suitable for every task, but can be quite successful when it allows users to explore and interact with the environment in order to accumulate information about the underlying data that is represented. In this case, exploration will help the user build a mental mapping between their visual system and those of other beings. The underlying data is actually used to build different perceptions of the scene.

Another advantage of using computer-generated environments as a means to showcase the simulations is that it lends well to incorporation into educational games. In order to appeal to younger users, in particular, it appears increasingly important to embed the learning material within an interactive, experiential application such as a game [14, 18, 22, 85]. Researchers have devoted considerable attention in attempts to harness the interest and engagement that drive people to play games so that it can be used to motivate educational material. These ideas are referred to as ‘game-based learning’, ‘serious games’, or ‘edutainment’.

Computer games have been the focus of several recent perceptual experiments. Giving the user a task, such as playing a game, can modify their fixation behavior [77] and their ability to notice level of detail changes [45]. It can also lead to inattentional blindness, where the player focuses so heavily on their goal that they lose sight of elements unrelated to it, which can compromise the educational impact of a game [8, 71]. The educational games developed to incorporate and enhance my simulations use the idea of task-specific perception as a means of emphasizing the information being presented.

Early attempts to insert learning into games received mixed results. The general approach taken was to present the educational material to the player and then to reward them for demonstrating their knowledge with a small quantity of gameplay. Some took the opposite approach of making the learning artifacts the reward for doing well at the game: some snippet of knowledge was revealed when the player accomplished a feat in the game [22]. In short, the learning was *extrinsic* from the game mechanics [31]. These approaches were rarely successful, as they treated the learning and the game as two separate elements [22]. The player’s desire to master the game had little connection to the learning material, and they came away remembering only the game, not the educational content. This is consistent with research in education that suggests employing new media, such as computer games, will only be effective if the delivery mechanism compliments the design of the content and the instructional methods [34].

Recent research in game-based learning has taken a more measured approach, recognizing that the key is to embed the learning content into the game in such a way that learning it is beneficial to the player’s success in the game—in other words, to make the educational content *intrinsic* to the game’s mechanics [31]. This idea draws from the principle of situated cognition [18], and is linked to experiential learning. Experiential learning suggests that retention can be improved if the knowledge in question is linked to a particular experience, increasing the chance that the knowledge will be integrated into the participant's cognitive processes [61].

Success has been achieved by integrating educational material into the early parts of the game design process. This is consistent with Gee's work [22], where he identifies 36 learning principles common within video games. One can take advantage of these principles by inserting the learning material into the right parts of the game. In addition to stimulating interest and engagement, this has shown particular potential for long-term retention of information [56]. The challenge then becomes largely one of effective game design. Previously, the metric used to judge the success of a particular game design was the amount of 'fun' it generated for the player. This is a fairly nebulous measure, that depends on ones' subjective experience and views. When considering educational games, one is essentially modifying the metric that is used to judge game design so that it includes a measure of how effectively the educational material is conveyed to the player. A successful educational game design is one that delivers a sense of 'fun' while also teaching the desired content [39, 53].

Evaluating the success of educational games is particularly challenging. The approach taken in this dissertation draws inspiration from previous work in evaluating video games in general, and educational games in particular. The most common method of evaluating the enjoyment or engagement of a game is through a post-test with Likert items dealing with various aspects of the experience. Several different instruments have been proposed in the research community. Fu et al. [20] proposed a model of engagement based on flow that evaluates the engagement in a game using a Likert scale with 42 items across 8 dimensions. O'Brien [57] proposed a more general model of engagement that is meant to evaluate tasks involving technology. This dissertation employs O'Brien's model as part of the evaluation for the educational games created based on the simulation framework, as described in Chapter 11.

Several evaluation approaches have also been documented that use physiological measures to determine the enjoyment derived from a game. Mandryk [49] used heart-rate and galvanic skin response to help determine whether a player was fully committed to a game. This approach can give better on-the-moment responses than a questionnaire, but is often difficult and costly to employ. Furthermore, it is not suited for testing the educational material conveyed by a serious game.

Evaluating educational games requires different testing instruments, as an emphasis must be placed on the information being conveyed by the game. A standard procedure is to use pre- and post- tests to determine how the players' knowledge has changed after playing the game [58,61,64,85]. Sometimes a control group is employed to help compensate for any knowledge that might be gained just by responding to the test instruments. Survey instruments are generally considered more effective when evaluating games intended to teach processes, such as mathematics, because the pre- and post- test instruments can easily be devised to test similar but not identical problems to the ones used in the game. Once developed, an instrument can be tested with a control group to determine if it has any effect on learning, and can be re-used to evaluate other games that teach the same process [39].

Naden [56] suggests that surveys are not as effective for testing fact-based knowledge, as their responses can be shallow and difficult to interpret. This makes it more challenging to test with a pre-/post-test experimental design, because the superficial results from a multiple choice test do not necessarily indicate that the respondent has achieved a deeper understanding of the subject matter. Naden suggests coupling survey results with game-play data such as response time or user mouse movements to get a more complete look at how the game is being played. In addition, the usual pre/post-test Likert evaluation structure can be augmented with more open-ended interviews to help determine the depth of knowledge that may have been obtained by the participants, and whether they are able to analyze and assess the knowledge, which would indicate that they have achieved a higher level on Bloom's taxonomy of learning [2].

Chapter 3

VDS Methodology

The VDS methodology is inspired by research in non-photorealistic rendering that uses factors of human perception to produce more compelling imagery [5, 24, 25]. Any imagery or rendering intended for human consumption will, by necessity, be passing through the channels of the human visual system, and this can be exploited in order to produce visuals that are more compelling to human observers.

3.1 Methodology

The aspect of human perception of greatest relevance to the VDS methodology is the human visual system's reliance on relative assessments, rather than absolute values [24]. The sensors in the human visual system are continually adapting to the levels of light reaching them, and this makes us more adept at sensing change, contrast, and edges, than at assessing absolute intensity values [43, 46]. NPR researchers have taken note of this characteristic when developing methods for tone-mapping [86], colorizing grayscale images [24, 25], and constructing visual illusions [9].

By drawing inspiration from this aspect of human perception, a powerful methodology can be developed for simulating visual systems in terms of how they relate to human vision. I propose a *Visual Differences Simulation* framework that portrays information about animal vision by highlighting the ways it differs from human vision. Note that this is not intended to provide a complete, physiologically accurate view of how an animal species sees the world—indeed, research into exactly how other species perceive the world is far from conclusive. Instead, the aim is to illustrate the main ways in which an animal's visual characteristics *differ* from our own. The average person does not know the processes at work in their visual system, but they are familiar with the way they see the world and can use their perception as a point reference when examining the simulations produced using the VDS methodology.

Based on the VDS methodology, the task at hand is one of identifying the major visual characteristics that can differ between visual systems and finding ways to map those differences into a rendering of a scene. There are three interacting factors that influence how one might choose to represent differences in each visual characteristic, described below:

- Semantics: It is critical that the methods of representing visual characteristics are as *meaningful* as possible, such that an unprimed observer will be able to glean some intuitive notion of how the visual system in question differs from human vision. This criteria considers whether the visual cue(s) chosen contain the appropriate semantics for conveying the visual characteristic being portrayed.
- Independence: Each representation in the simulation must be sufficiently independent from the others so as to avoid interfering with them. For example, relying too much on a single visual cue (such as color) can overload it, and consequently reduce its discriminatory power.
- Efficiency: The chosen representation must be sufficiently efficient that it can be executed in real time on commodity graphics hardware. A real time simulation is important for creating the interactive and immersive experience of looking out through another's eyes, and so algorithms and methods must be chosen that do not require sacrificing performance.

The following six chapters discuss how the VDS framework represents differences in visual characteristic using the difference-based methodology, and how the three factors listed above inform the selection of the most appropriate visual representation. The description of each characteristic includes several examples where the representation is used to simulate an aspect of an animal's visual system. Figures 3.1, 3.2, and 3.3 show how the transformations are used together to compose the illustrative simulations for the cat, pit viper, and bee visual systems. The parameters used by the VDS framework to produce these simulations can be found in Appendix A. Figure 3.4 shows a small sample of the visual systems that can be represented using the VDS framework.

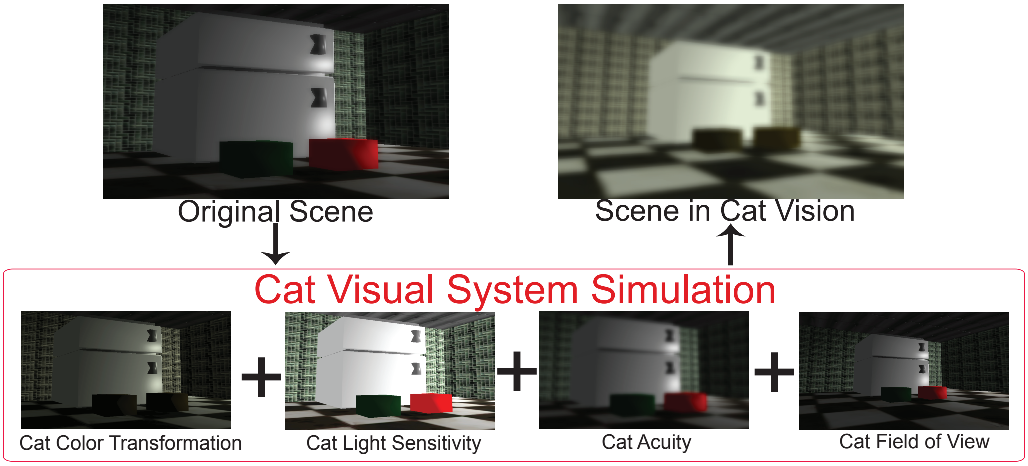

Figure 3.1: The cat visual system simulation includes four transformations intended to illustrate four major differences between the human and cat visual systems.

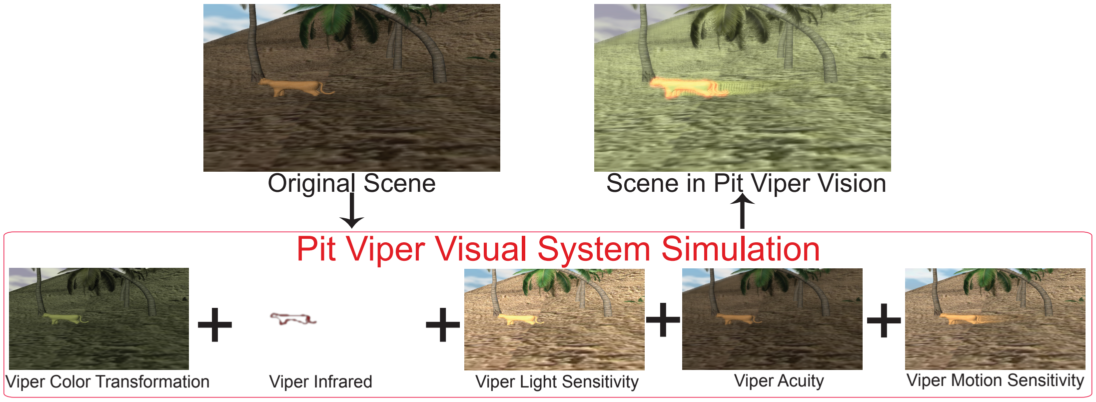

Figure 3.2: The pit viper simulation illustrates five of the major differences between the human and pit viper visual systems, including color discrimination, visual acuity, hypserspectral sensitivity, motion tracking, and the minimum light detection threshold. These transformations give the user a sense of how their vision differs from that of a pit viper.

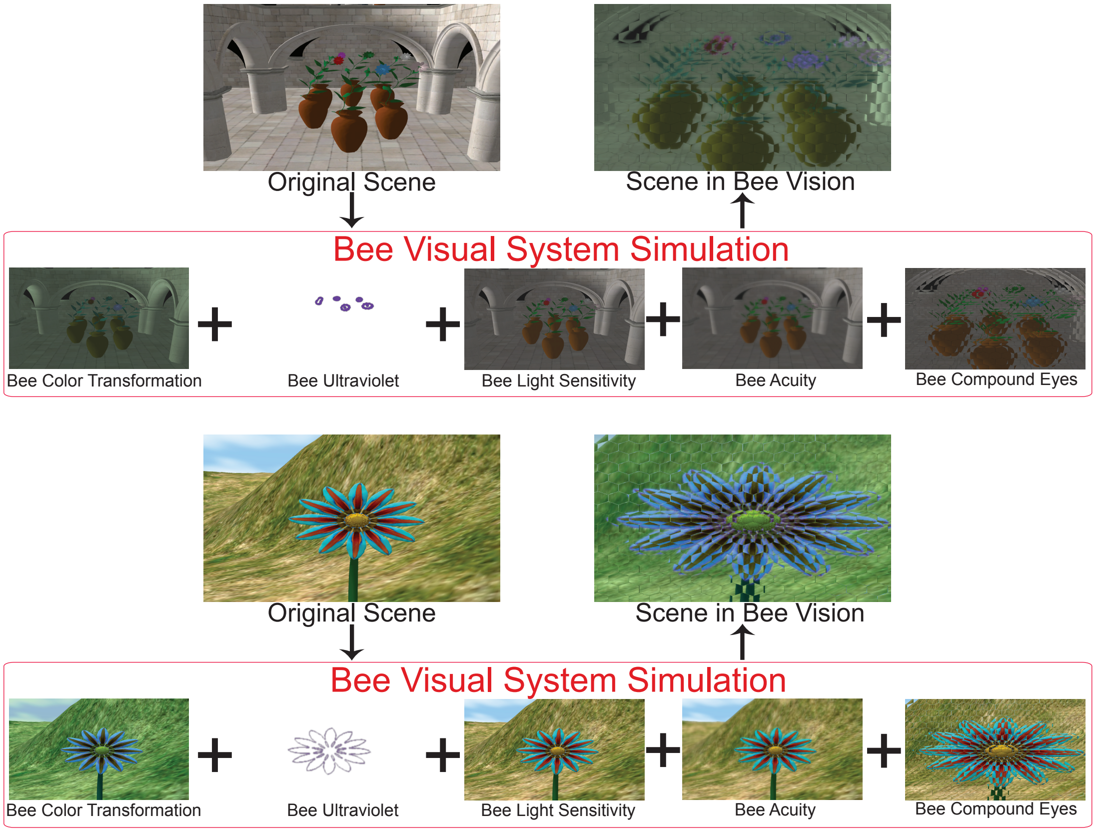

Figure 3.3: The bee simulation represents five of the major differences between the human and bee visual systems. Four of these transformations are performed primarily by fragment shaders, allowing them to be executed in real time on a virtual environment or a live video feed, provided the inputs to the simulation are properly specified.

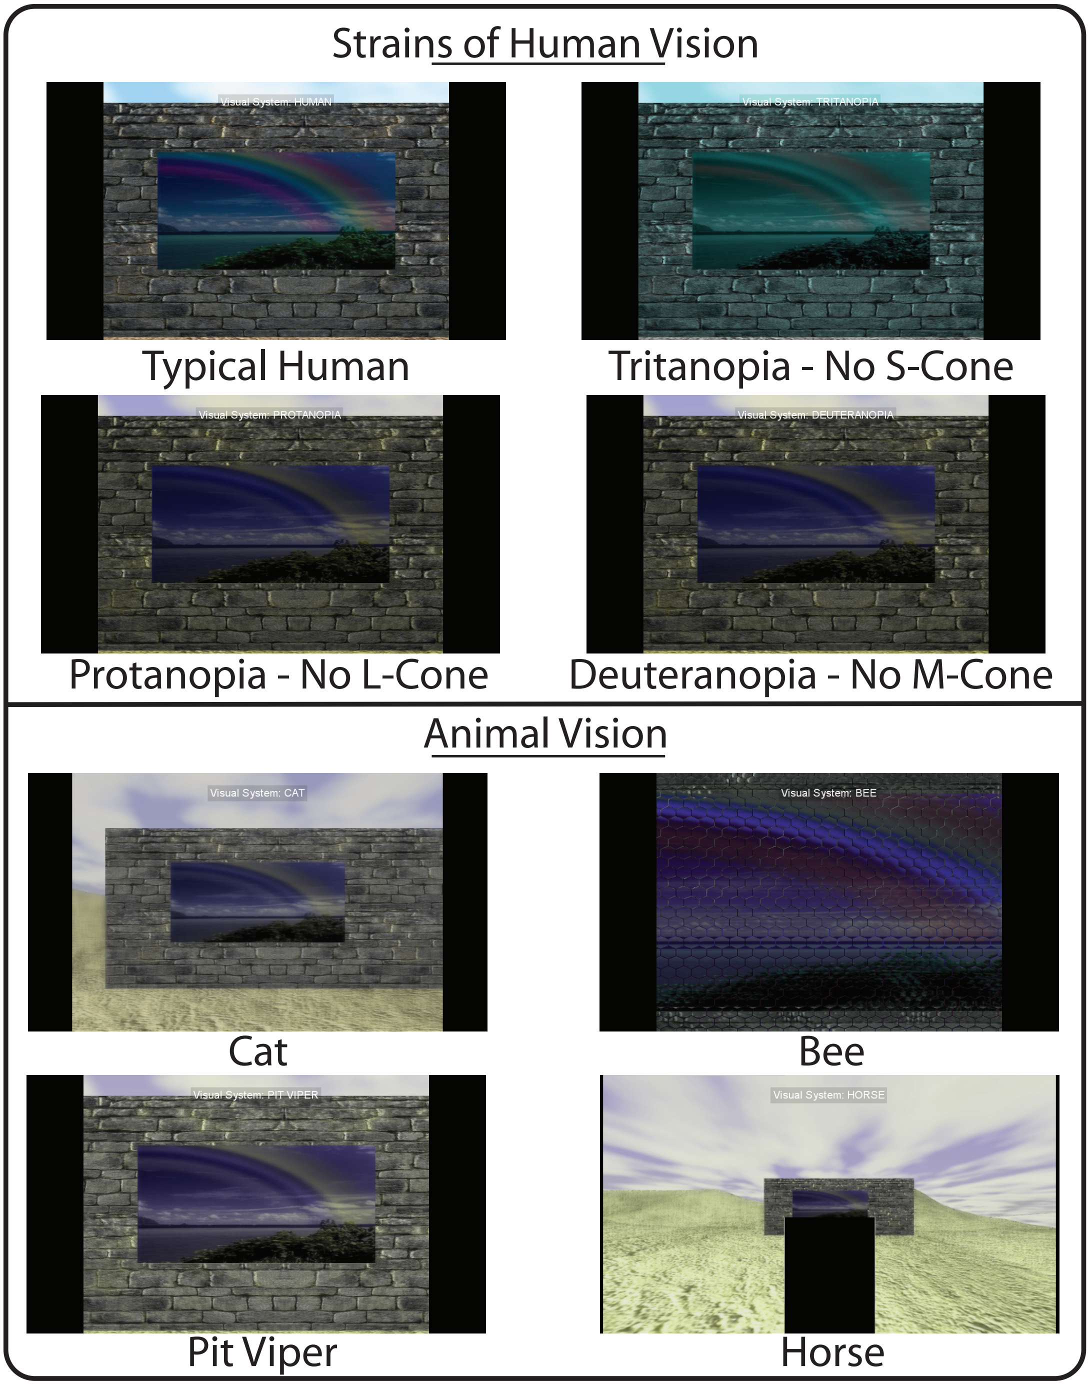

Figure 3.4: The VDS framework is capable of representing several different strains of human vision, including dichromancy strains such as tritanopia, protanopia, and deuteranopia. The top left image shows the scene as it would be observed by the typical human visual system, and serves as the input image that can be transformed by the framework to portray the other types of human vision. The framework is also sufficiently general to illustrate a variety of animal visual systems. The transformations that make up the simulations run in real time on a virtual environment, and this allows for the observer to learn about and experience these visual systems in an interactive and exploratory setting.

Chapter 4

Color Vision

Cones are the primary photoreceptors associated with color vision, and contain pigment that is sensitive to different wavelengths of light [84]. Many visual systems also include rod photoreceptors, which mediate vision under low light conditions. The interaction between rods and cones under low light (scotopic or mesopic) conditions is complex, and beyond the scope of the color transformation described here. This model of color vision assumes that the scene is seen under photopic conditions, where the cone photoreceptors are dominant.

The human visual system typically includes three types of cones that respond to different parts of the light spectrum. The colors that we perceive are determined by the relative amount that a particular waveform of light stimulates the three cone classes. The left part of Figure 4.1 shows the spectral sensitivity curves of the human cone cells, as represented by Stockman et al.'s cone fundamentals [74, 75]. It is important to note that 'colors' are descriptors for the way humans perceive certain wavelengths of light, but are likely not interpreted the same way by animals when looking at similar waveforms of light [43]. This makes it difficult to devise a physiologically accurate means of simulating the color sensitivity of different visual systems.

Instead of attempting to simulate color vision in an absolute manner, the VDS methodology suggests that color discrimination be represented *relative* to typical human color vision. The color transformation can then be used to illustrate the parts of the human spectrum that the system being simulated can detect, even if they may not *perceive* those wavelengths of light the same way that humans do. Consequently, the two goals of the color vision transformation are as follows:

- Demonstrate the parts of the human spectrum that would **not** be visible to the

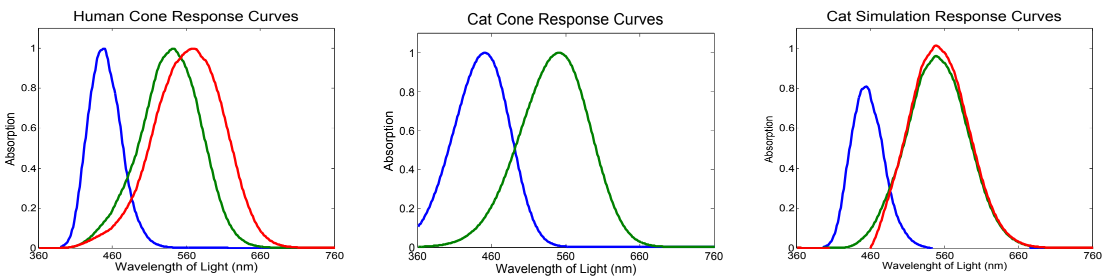

Figure 4.1: Response curves for cone photoreceptors in the human visual system (left), the cat visual system (center) and the human response curves after they have been modified to remove light that would be invisible to cats (right).

system being simulated;

- Demonstrate where the color spectrum of the species being simulated lies relative to the typical human spectrum.

In order for the colors shown in the simulation to be meaningful to an unprimed observer, they should be kept as consistent as possible with their human interpretations. For example, red shown in the simulation should represent wavelengths of light around the range of the spectrum that humans perceive as red, rather than treating wavelengths of light that stimulate the animal's third cone class as the color red, regardless of what color that light would appear to humans.

Keeping colors consistent with their human interpretation allows observers to establish a mental mapping between the wavelengths of light that make up the human spectrum and the spectrum of the visual system being simulated. This should help make it apparent where the animal's color range falls relative to the typical human spectrum, even if it is not known how the animal would actually *perceive* those wavelengths of light. When the animal's perception looks more blue, it is because they are more sensitive to shorter wavelengths of light that humans perceive as blue or purple. In a sense, this creates a perceptual anchor on which to ground the observation of the colors in the simulation.

The algorithm used by the VDS framework for simulating differences in color vision begins by generating a set of spectral sensitivity curves for the visual system being simulated. Each type of cone photoreceptor contains a slightly different pigment that determines the spectrum of light to which it is sensitive. Some researchers have suggested that this pigment has relatively invariant properties across different species. This notion is referred to as a universal pigment template, and it simplifies analysis by allowing the spectral response curve of a cone receptor to be defined

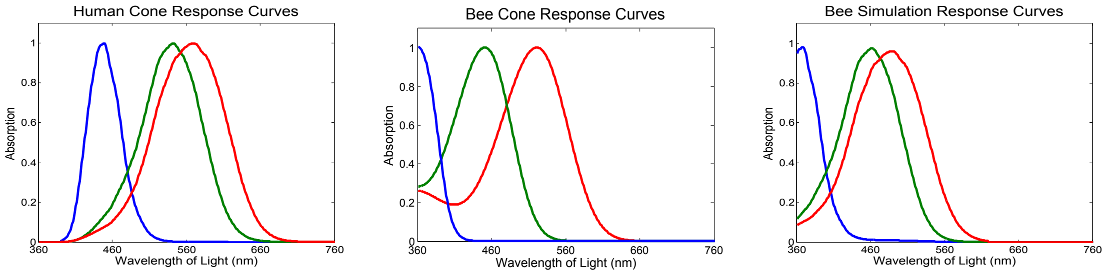

Figure 4.2: Response curves for cone photoreceptors in the human visual system (left), the bee visual system (center), and the human response curves after they have been shifted and modified to remove light that would be invisible to bees (right).

by just specifying its peak absorption point. Several different formulations of the universal pigment template have been proposed that take different data sets and visual systems into consideration [40]. The VDS framework uses the alpha-bands proposed by Govardovskii et al. [28] to create response curves based on peak absorption data for the system being simulated. Govardovskii et al.'s template is based on bovine data, but the sacrificed precision is unlikely to be noticed by observers of the simulation. Figures 4.1 and 4.2 show response curves generated for cat and bee visual systems using this universal pigment template.

Given discretized response curves for humans and some other visual system, the goal is to devise a new set of cone response curves that contain only radiance distributions that are visible to both humans and the other visual system that will be simulated. Appendix B describes the linear algebra used to derive a matrix S that can transform the colors in the scene so as to remove those that would be invisible to the system being simulated.

The method in Appendix B assumes that the wavelengths of light perceived by the system being simulated are approximately a subset of the human spectrum. In order to generalize the method, one needs to consider the case where the system being simulated is sensitive to different wavelengths of light from the human cone photoreceptors, such as the bee spectral sensitivity curves shown in Figure 4.3. As such, the color transformation starts by shifting the human spectral sensitivity curves such that the low wavelength cone peak is aligned with the lowest wavelength cone peak of the visual system being simulated before running the algorithm in Appendix B to generate color matrix S.

The alignment process is necessary to allow the color transformation matrix to be generated for any arbitrary set of spectral sensitivity curves. The alignment needs to be represented visually so as to illustrate that the visual system being simulated

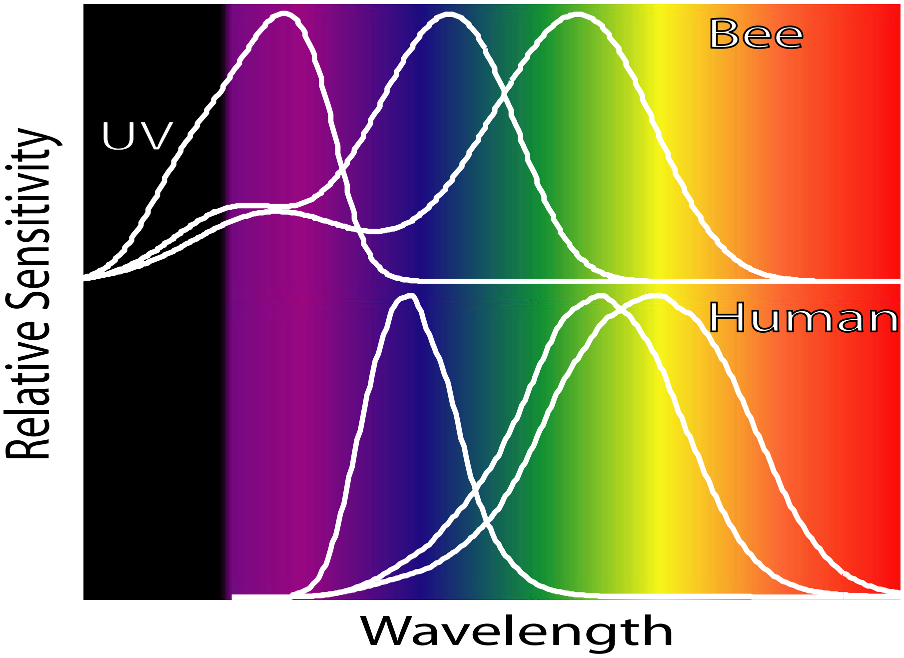

Figure 4.3: Approximate spectral sensitivity curves for the bee color vision system, based on spectral sensitivities measured by Menzel et al. [51], and Stockman et al.'s human cone fundamentals [74].

sees a spectrum that is shifted from the typical human spectrum. This is achieved by rotating the hues in the scene such that the color palette more closely represents where the visual system's sensitivities lie relative to the human spectrum.

The first step in the hue shift process is to convert the colors in the scene into the HSL color space. The resulting hues are then rotated by the same amount that the human curves are shifted to generate S. In order to be consistent with human color perception, the hue rotation is clamped at the 330 degree mark (shown in Figure 4.4). The rotation through hue space effectively illustrates where an animal's color spectrum falls relative to the human spectrum. For example, simulating animal visual systems sensitive to lower wavelengths of light requires shifting the human curves in the negative direction, which leads to a negative rotation through hue space that makes the scene appear more blue and green, colors associated with lower wavelengths of light. Conversely, animals sensitive to higher wavelengths of light see more wavelengths that humans associate with orange and red, leading to a positive hue rotation that produces more orange and red colors. The hue rotation runs in real time as part of the same fragment shader that performs the color transformation.

The saturation of the colors in the scene can also be adjusted during the color transformation to represet species with less sensitive color perception. Saturation

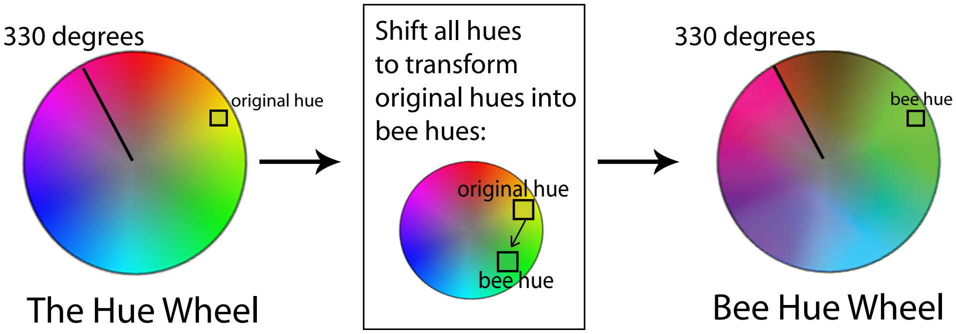

Figure 4.4: The color transformation rotates colors around the hue wheel in order to show where the color space of the system being simulated lies relative to the human spectrum. The rotation is clamped at the 330 degree mark in order to be consistent with human color perception. For example, the bee color space is represented by rotating each hue clockwise in order to create a set of hues that illustrate bee sensitivity to the range of the spectrum that humans perceive as blues and greens.

changes are achieved by scaling the S channel of HSL values in the scene. Reducing the S values results in colors that appear faded and desaturated, which is effective for representing visual systems such as the pit viper.

Some species have more than three types of color cones, allowing them to disambiguate between colors that look the same to humans. This is difficult to represent in a three-dimensional color space like RGB. One method to achieve this would be to come up with a new space with n primaries, where n is the number of cone classes in the visual system, but this would be difficult to reconcile with the VDS methodology that attempts to keep the colors consistent with how they would be perceived by the human visual system.

4.1 Evaluating the Color Vision Transformation

4.1.1 Semantics

A difference in color discrimination should have a meaningful effect on the colors in the scene. The VDS framework transforms the colors in a manner that makes clear where the animal's color discrimination falls relative to the human spectrum. In other words, the colors shown in the scene bear relation to the way humans would perceive those wavelengths of light. For example, a simulated system that sees more lower wavelengths of light is presented using a color palette consisting primarily of colors that humans associate with lower wavelengths, such as blues and greens. Light from outside the human spectrum is handled separately with the transformation described in Chapter 5.

4.1.2 Independence

Color is one of the most noticeable visual cues in a scene, and has a tendency to get overloaded with significance [24]. As such, care needs to be taken to ensure that its discriminatory power is preserved. In this case, the color vision transformation shows a difference in color discrimination, making changes to the displayed colors warranted.

4.1.3 Efficiency

The color transformation used by the VDS framework requires calculating a color transformation matrix as a pre-processing step that represents the differences between the human color cone sensitivity curves and the sensitivity curves of the species in question. At run-time, a fragment shader is used to multiply the colors at each pixel by the color transformation matrix. The results of this multiplication are then converted into the HSL color space. The hues are rotated, if necessary, and the result is converted back into RGB for display on the screen. These operations run in real time as part of a fragment shader.

4.2 Color Vision Examples

4.2.1 Cat Color Vision

Schuurmans and Zrenner [68] reported that the cat visual system includes two cones classes, with the short wavelength cones peaking at 450 nm, and the medium wavelength cones peaking around 556 nm. This model of vision implies that cats would only be able to see combinations of two color primaries. Some researchers have suggested the presence of a third cone class within the cat visual system [29], but they believe this cone class would only be present in a very small number of retinal ganglion cells, limiting its ability to mediate color vision. It is difficult to predict the exact influence such a cone class would have on color vision, but it would present the possibility that cats could, perhaps with difficulty, distinguish between colors on the higher end of the human spectrum. I choose to adopt the two cone model of the cat visual system, as it is well suited to illustrate the difficulties cats may have in distinguishing the higher wavelengths of light that humans perceive as greens and reds.

The modified human response curves for approximating cat color vision are shown in the third image of Figure 4.1. The color transformation is produced using a single 3x3 matrix multiply, which is applied in real time using a fragment shader. Figure 4.5 shows images of a computer-generated environment before and after the application of the cat color transformation.

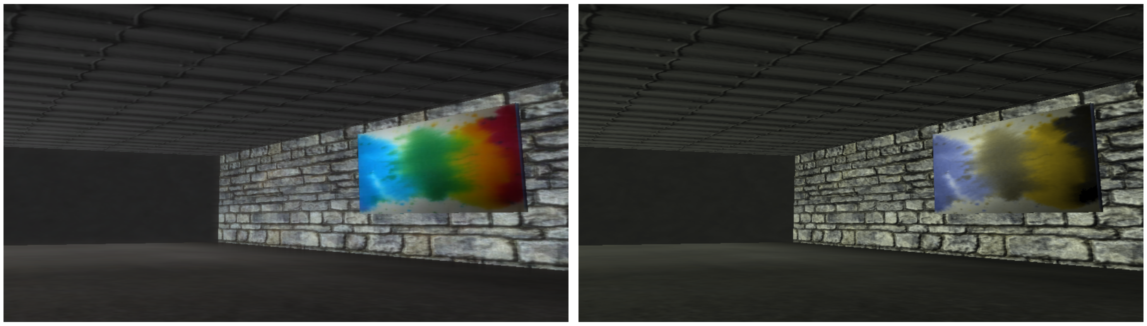

Figure 4.5: The source image (left) and the result of the cat color transformation (right). Cats are thought to have difficulty discriminating higher wavelength of light that humans perceive as oranges and reds. The simulation effectively removes these colors from the scene, while otherwise keeping the colors consistent with the wavelengths with which they are associated by the human visual system.

4.2.2 Bee Color Vision

Spectral sensitivity curves for the bee color vision system are generated using Govardovskii et al.'s universal pigment template [28] and the bee cone peak absorbance points measured by Menzel et al. [51]. Figure 4.2 shows a comparison of the human and bee spectral sensitivity curves. As described above, the human curves are shifted down in wavelength in order to align with the low wavelength cone peak of the bee visual system. This shift is accompanied by a rotation through hue space that emphasizes the blue and green colors in the scene. Figure 4.6 shows the result of performing the bee color transformation on an input scene.

4.2.3 Pit Viper Color Vision

Figure 4.7 shows the result of the pit viper color transformation. It is believed that pit vipers have two cone classes that peak around 430 and 550 nm. This gives pit vipers sensitivity similar to humans who are red-green color blind. Pit vipers are also

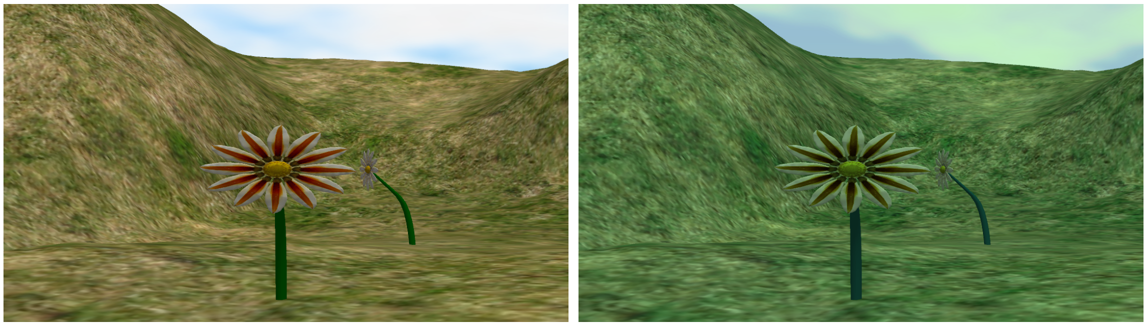

Figure 4.6: The source image (left) and the result of the bee color transformation (right). Bees perceive lower wavelengths of light than humans. The transformation represents this by showing more blue and green colors that humans associate with lower wavelengths of light, and the reduction of colors that humans associate with higher wavelengths of light, such as reds and oranges.

able to detect infrared light from outside the human spectrum that feeds into their visual channel, but that is represented by the hyperspectral sensitivity transformation described in Chapter 5.

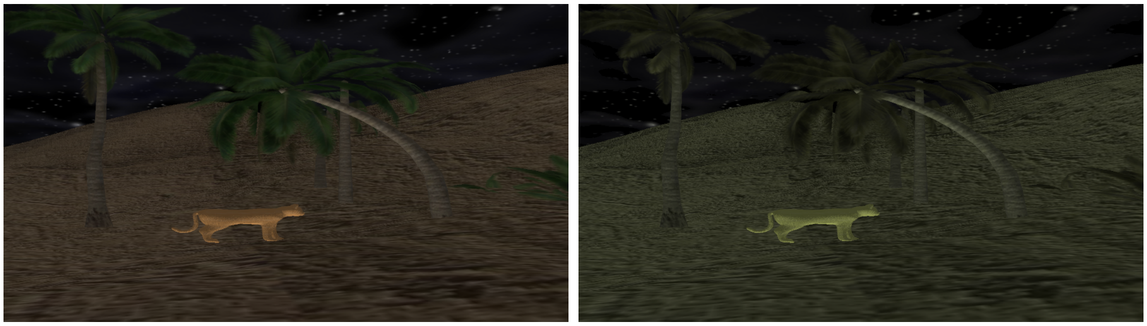

Figure 4.7: The source image (left) and the result of the pit viper color transformation (right). Pit vipers are thought to perceive wavelengths on the edges of the human spectrum, allowing them to detect hyperspectral light such as ultraviolet and infrared wavelengths, but they are not able to discriminate across the human range as effectively. The simulation gives the user this impression by showing a desaturated view with a limited color palette.

Chapter 5

Hyperspectral Sensitivity

Representing hyperspectral sensitivity is difficult, because hyperspectral light has no clear visual appearance within the domain of human perception, and yet a way must be found to map it into one of the available visual cues. There are thus three separate goals to be satisfied:

- Represent ultraviolet and infrared light in a meaningful manner;
- Avoid detracting from the discriminatory power of the other colors in the scene, which are already meaningful because of the transformation described in Chapter 4;
- Indicate that hyperspectral data is outside the human spectrum, and is thus visually ambiguous from a human perspective.

A common approach for representing hyperspectral data is through false coloring, where the hyperspectral data is visualized using a colormap. This does not make it clear that the hyperspectral data is visually ambiguous, except insofar as the colormaps chosen when using this approach tend to be unnatural. By overloading the visual cues provided by the colors in the scene, an unprimed observer could potentially confuse this effect with the color transformation described in Chapter 4, compromising the semantics of both effects, and violating all three of the goals listed above.

Instead, the VDS framework represents hyperspectral sensitivity using a dynamic glow effect that incorporates transparency. Transparency has been noted as an effective way for conveying uncertainty, although humans have limited discriminatory power with regard to differences in transparency values [87]. For illustrative purposes, an almost binary form of hyperspectral sensitivity is acceptable. Furthermore, the

VDS framework can take advantage of its interactive framerate by using time as another variable in depicting the presence of hyperspectral data. The transparency in the hyperspectral transformation varies dynamically over time, drawing the attention of the human eye, while reducing the impact on the color transformation described in Chapter 4.

The second constraint can also be supported by annotating only the silhouette edges of an object with the glow effect, further reducing the changes made to the colors produced by Chapter 4. This is particularly appropriate when simulating systems like bee vision, where it is important to indicate that some flowers reflect ultraviolet light, and that this helps bees locate and identify them. However, most flowers do not reflect ultraviolet light uniformly. Instead, they have patterns of reflective pigment spread across their petals. Using the glow effect to highlight the silhouette edges helps to represent this. A more complex solution would be to store a separate set of textures for each model that specify the pattern of ultraviolet light that it reflects, but such a solution would put a greater burden on asset creation, and thus would put more constraints on the inputs necessary to run the simulation.

The glow effect is achieved by first rendering a mask of all the visible objects that are set in software to reflect hyperspectral light. Next, a fragment shader finds the silhouette edges of these objects, and alpha-blends the edge colors with the color specified in software to represent hyperspectral light—purple for ultraviolet (RGB: 0.5, 0, 1.0), red for infrared (RGB: 1.0, 0, 0)—multiplied by the hyperspectral surface reflectance properties, which can be specified per object and must fall within a 0 - 1.0 range. The transparency levels used for the alpha-blending vary between 0.25 to 1.0 using a sinusoidal function based on the number of milliseconds elapsed. Finally, the edges are subjected to a Gaussian blur to complete the effect. Figure 5.1 shows an input image and the results of the glow effect, before it is alpha-blended with the output of the other transformations that make up the simulation.

Section 12.3 describes an experiment that was conducted to evaluate the effectiveness of the glow effect relative to a more traditional false coloring approach. Users were shown flowers of different colors from within the bee color space, and told to determine the colors they would correspond to in the human color space. Some of the flowers were specified to reflect ultraviolet light, and in the first condition the ultraviolet contribution was represented using false coloring with a purple colormap, whereas in the second condition the ultraviolet light was represented using the glow effect described above.

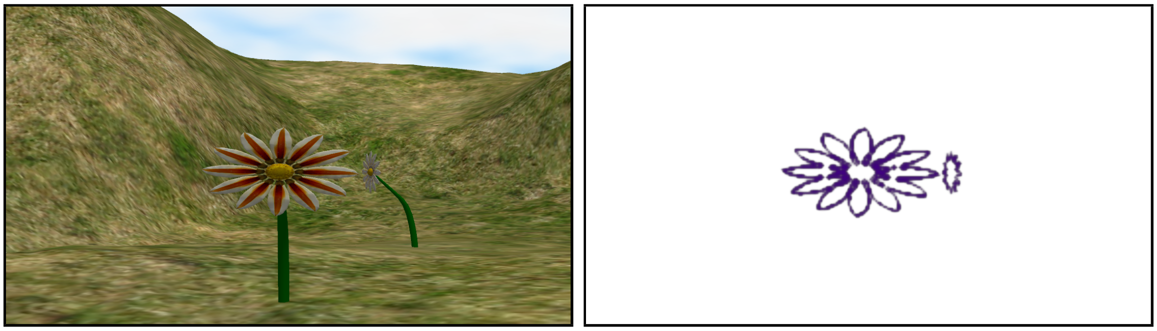

Figure 5.1: The source image (left) and the result of the glow produced by the hyperspectral visualization effect, before being composited with the results of the other transformations that make up the bee simulation (right).

The experiment determined that the number of false negatives increased significantly in the false coloring condition as the contribution of ultraviolet light increased. In other words, the respondents failed to identify the flower colors more often as the intensity of the false coloring effect increased. However, there was no significant increase in false negatives when the glow effect was used. This could potentially mean that the glow effect has less impact on the other colors in the scene, and is thus more independent from the color transformation. This experiment is described in more detail along with the rest of the evaluation in Section 12.3.

5.1 Evaluating the Hyperspectral Sensitivity Transformation

5.1.1 Semantics

The VDS framework represents hyperspectral data using a glow effect that combines transparency, dynamics, and false coloring. Including color as part of the equation seems justified given that this effect is representing wavelengths of light that are visible to the visual system being simulated. However, it would be misleading to use color alone, because that would exclude the important semantic information that the colors are representing wavelengths of light from outside the human spectrum. The framework addresses this by using transparency as part of the glow effect, and that has been shown to be a successful way of representing uncertainty [87]. This conveys that the hyperspectral colors are visually ambiguous from a human perspective.

5.1.2 Independence

The hyperspectral representation involves changing some of the colors in the scene. This contributes towards overloading the color cue, reducing its discriminatory power and potentially compromising the effects of the color transformation. However, this is mitigated by incorporating transparency, and making the effect dynamic. The influence of the hyperspectral representation oscillates over time, so there will be moments when it has a negligible effect on the colors in the scene.

5.1.3 Efficiency

The hyperspectral representation requires a separate render pass to produce the mask of elements that are marked in software to reflect hyperspectral wavelengths of light. This mask is passed into a fragment shader that determines the silhouette edges of the objects subject to the effect, and then the fragment shader highlights the edges with the hyperspectral colors. In spite of the extra render pass, this effect can still be employed in real time.

5.2 Hyperspectral Sensitivity Examples

5.2.1 Bee Hyperspectral Sensitivity

Bees are capable of perceiving ultraviolet light that is invisible to human eyes [16]. Many flowers have patterns of pigment that reflect ultraviolet light across their petals, allowing bees to identify and pollinate them. Figure 5.1 shows the result of computing the hyperspectral transformation, before it is composited with the results of the other transformations that make up the bee simulation. The transparency of the effect varies over time, indicating that the ultraviolet light has an uncertain appearance within the human domain.

5.2.2 Pit Viper Hyperspectral Sensitivity

Pit vipers are capable of detecting infrared wavelengths of light thanks to heat-sensitive pit organs located near their eyes. It is not entirely clear how this data is integrated into their visual perception, but it is believed that pit vipers use this hypserspectral data to help them catch prey. The signals detected by the viper's pits are thought to be quite blurry, as point-like objects spread into disc-shaped images on the pit membranes [69]. The glow effect used by the VDS framework includes a

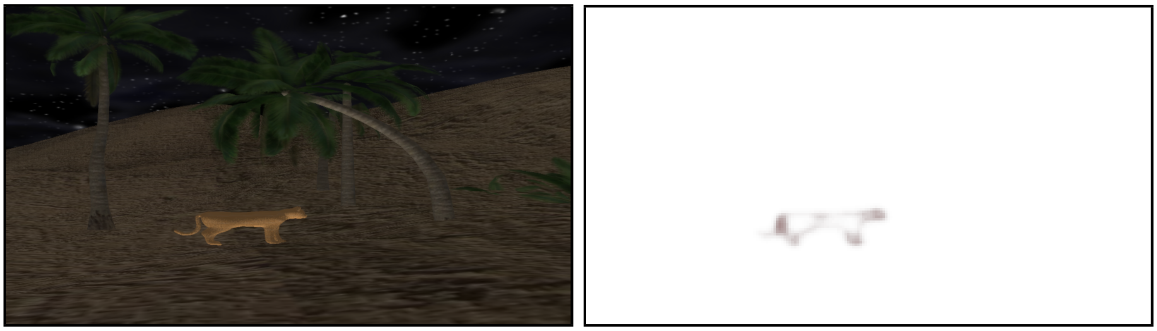

Figure 5.2: The source image (left) and the glow produced by the hyperspectral representation, which will subsequently be composited with the results of the other transformations that make up the pit viper simulation (right). Pit vipers have membranes near their eyes that are sensitive to infrared wavelengths of light that help them track down living prey. This is visualized using a glow effect with a red hue to show that the infrared wavelengths of light are visual ambiguous from a human perspective.

blur component, making it particularly suitable for representing this characteristic. Figure 5.2 shows the result of computing the hyperspectral transformation for the pit viper visual system, before it is composited with the results of the other transformations that make up the pit viper simulation.

Chapter 6

Visual Acuity

Visual acuity is the sharpness of the perception produced by the visual system. Acuity is determined by the shape, construction, and placement of the eyes. Human eyes, for instance, have a concentrated region of cone photoreceptors called the *foveal region* that is effective for focusing on specific areas. It is thought to have evolved as an asset for discriminating between edible and poisonous plants while foraging for food [43]. The peripheral regions of the human retina have much sparser cone coverage, and as a result produce more blurred views [60].

Many animals do not possess a dense foveal region, as their evolutionary path has not required them to focus on fine details around them. Instead, species such as cats and horses have a concentrated band of photoreceptors called a 'visual streak' that is more sensitive to detecting movement along the horizontal plane [55, 65, 68]. The VDS framework takes motion tracking mechanisms into account in Chapter 8.

Differences in visual acuity can be simulated using a blur, and can be incorporated into depth of field effects that are common in real time rendering engines. Near and far acuity can be specifically represented using a depth blur fragment shader—a shader that uses a Gaussian blur with influence that varies according to distance from the camera. This requires a separate render pass that generates a depth map indicating how far each pixel is from the camera. The depth map is then used to influence a post-process blur on the rendered scene. The result is an image that grows more blurred as the distance from the camera increases, as shown in Figure 6.1.

The VDS framework applies a moderate distance blur to represent typical human vision, and either strengthens it or weakens it depending on the relative acuity of the system being simulated. It is worth noting that the magnitude of the blur need only be consistent relative to the amount of blur applied to the baseline human view. The *difference-based simulation* methodology is focused on representing the difference in relative terms, not in absolute magnitude. This is particularly apt in this case, because humans are already poor at judging distances in virtual environments [81], and thus would be unlikely to appreciate absolute precision even if it were feasible to present.

Equation 6.1 shows how the blur influence is calculated. The blur influence b is calculated based on c, the visual acuity of the system being simulated, specified in decimal notation. a is a constant representing the blur influence used in the basic human view. a is divided by c, adjusting the blur influence to represent differences in visual acuity. The images in this dissertation were produced with an a value of 0.25 which seems sufficient to make the blur noticeable but not overpowering. Depending on c, b will range from [0 - 1.0], and can be further modified based on the depth map.

$$b = a/c \tag{6.1}$$

6.1 Evaluating the Visual Acuity Transformation

6.1.1 Semantics

The semantics for the visual acuity transformation are intuitive—differences in acuity lead to differences in the sharpness of the resulting visual perception, so using a Gaussian blur whose influence varies based on the visual system's acuity relative to human acuity is a natural choice. The only potential problem is that small differences in blur are difficult to detect, and may go unobserved if they are not reinforced in some manner.

6.1.2 Independence

The visual acuity transformation blurs the colors in the scene, but its most pronounced effect is on the sharpness of the edges, and that is an intuitive consequence of differences in visual acuity.

6.1.3 Efficiency

Depth blurs have long been used in 3D rendering to achieve depth of field effects, and can be implemented in fragment shaders that run in real time.

6.2 Visual Acuity Examples

6.2.1 Cat Visual Acuity

It is estimated cats have visual acuity that is between four to ten times worse than humans. In medical terms, cats are thought to have 20/80 vision, meaning that what a normally-sighted human can see well at 80 feet, a cat can only see in as much detail at 20 feet [33]. Figure 6.1 shows the result of applying the cat visual acuity transformation, where the cat's view is considerably more blurred than the scene observed by the human visual system.

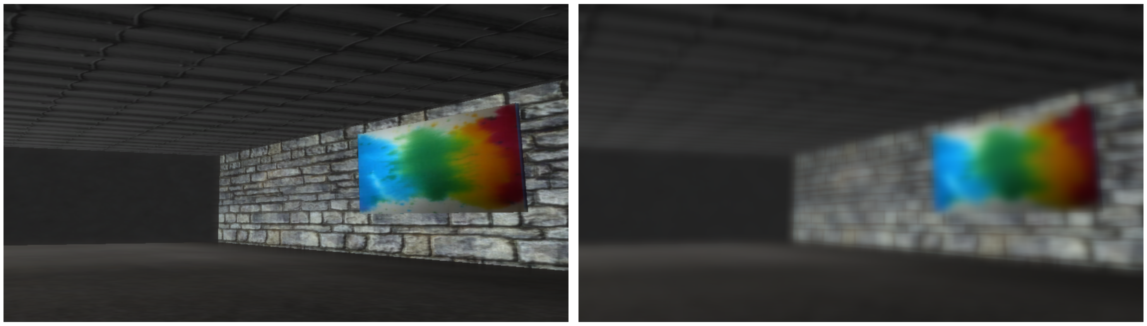

Figure 6.1: The source image (left) and the result of the cat acuity transformation (right). A blur effect is employed to represent that cats have considerably lower visual acuity than humans.

6.2.2 Bee Visual Acuity

Bees have far fewer photoreceptors in their eyes than humans; as a result, their visual acuity is worse. They also do not possess the dense concentration of photoreceptors in the foveal region that allows humans to effectively focus on specific points. Bee acuity is represented with a blur effect, where the strength of the Gaussian blur is adjusted based on the distance from the camera. The simulation also increases the strength of the Gaussian blur when the light intensity in the scene is low, because bees use neural summation to help navigate under low light conditions, and this serves to reduce their net visual acuity [10]. Figure 6.2 shows the result of applying the bee visual acuity transformation.

6.2.3 Pit Viper Visual Acuity

Figure 6.3 shows the result of the pit viper visual acuity transformation. Pit viper vision is optimized for light sensitivity and motion detection, and not for sharp reso-

Figure 6.2: The source image (left) and the result of the bee visual acuity transformation (right). Bees have fewer photoreceptors in their eyes than humans, and as a result their visual acuity is lower. The simulation represents this difference using a blur effect.

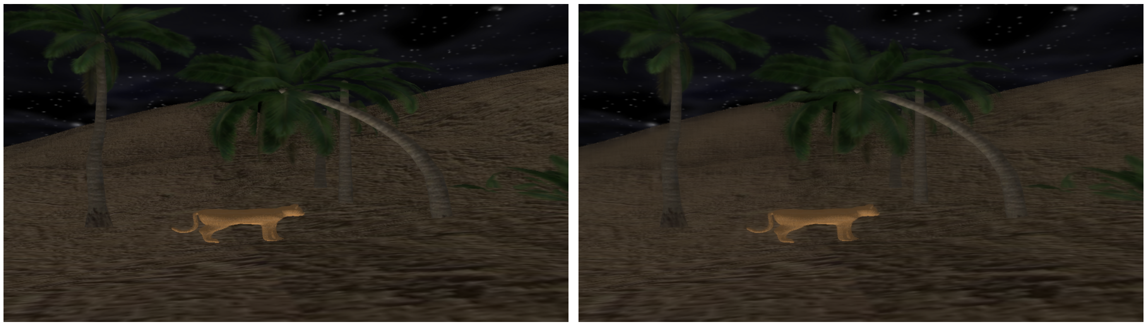

Figure 6.3: The source image (left) and the result of the pit viper visual acuity transformation (right). Pit vipers have inferior visual acuity compared to humans, and this is represented by applying a blur effect.

lution. Consequently, pit vipers have blurrier vision than humans. It is thought that they use the input from their heat-sensitive pit organs to help compensate for this deficiency and detect their prey.

Chapter 7

Light Sensitivity

The minimum light detection threshold is the point at which there is sufficient light to stimulate the photoreceptors in a visual system, and it is important for determining if the visual system can function under low light levels. As with the other characteristics, the goal of the VDS framework is not to model the absolute light sensitivity of a visual system, but instead to portray the visual system's light sensitivity *relative* to human light sensitivity.

The light sensitivity transformation in the VDS framework makes the assumption that the intensities of the lights in a virtual environment represent the light detected by a human observer. In order to represent the light sensitivity of another visual system, the framework employs a transformation that encapsulates their sensitivity relative to human light sensitivity. The framework supports two sorts of relations: linear and logarithmic. A linear relationship is one where there is a constant shift between the light detected by the human and the visual system being simulated. A logarithmic transformation is appropriate for representing species like cats that have lower minimum light detection thresholds, allowing their visual system to function in dark conditions. However, this enhanced light sensitivity causes the cat's photoreceptors to saturate at lower light levels, leading to a plateau in light sensitivity when light is abundant. Figure 7.1 shows an example of the logarithmic light transformation that the VDS framework employs to represent cat light sensitivity.

Both of the light sensitivity transformation functions are built into fragment shaders that can run in real time. Making these adjustments is an easy feat when the simulation is being run in a computer-generated environment, where the intensities of all the lights are specified in software, but considerably more difficult to accomplish if the simulation is run on a video feed, for instance, where the intensities of the lights are not explicitly specified.

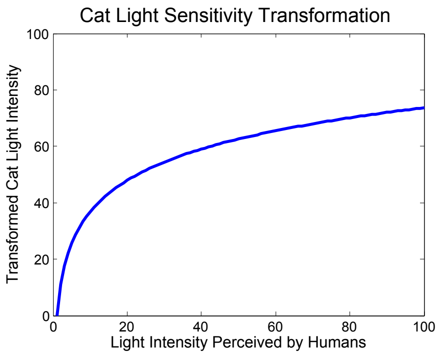

Figure 7.1: The result of the cat light sensitivity transformation graphed relative to the light perceived by the human visual system. This logarithmic relation allows cats to function at lower light levels, but their receptors saturate when light is abundant, leading to a plateau.

7.1 Evaluating the Light Sensitivity Transformation

7.1.1 Semantics

The minimum light detection threshold is an important characteristic of a visual system with regard to determining its ability to function in low light environments. The VDS framework represents differences in light sensitivity by modifying the intensities of the lights in the scene, giving a sense of how this characteristic affects a species' ability to function as the light level varies.

7.1.2 Independence

The light sensitivity transformation impacts the colors in the scene, but humans are fairly adept at separating brightness from colors when observing stimuli [43].

7.1.3 Efficiency

This effect can be run efficiently as part of a fragment shader, provided the intensities of the lights in the scene are all explicitly specified. Otherwise, calculating the contributions of lights to the scene can be a challenging task. As it stands, this effect is able to run in real time on a computer-generated environment.

7.2 Light Sensitivity Examples

7.2.1 Cat Light Sensitivity Transformation

Cats have a lower light detection threshold than humans, allowing them to see more clearly at night. This is partially due to an abundance of rod sensors, which dominate the visual system under scotopic (dark) conditions. In addition, cats, like dogs and many other animals, have a *tapetum lucidum*, which is a reflective layer behind the retina that reflects light that passes through the retina back into the eye [6, 43]. Although this enhances vision under low light conditions, it reduces vision in situations where light is abundant. The characteristics described above give the cat a minimum light detection threshold up to seven times lower than that of humans, but a lower light saturation point.

The relationship between human and cat light detection is approximated using a logarithmic function. Figure 7.1 shows that the slope of the resulting curve approximates the cat's ability to function at low light levels, while also representing the manner that the light receptor responses plateau under conditions where light is more abundant. Figure 7.2 shows the result of the cat luminance transformation.

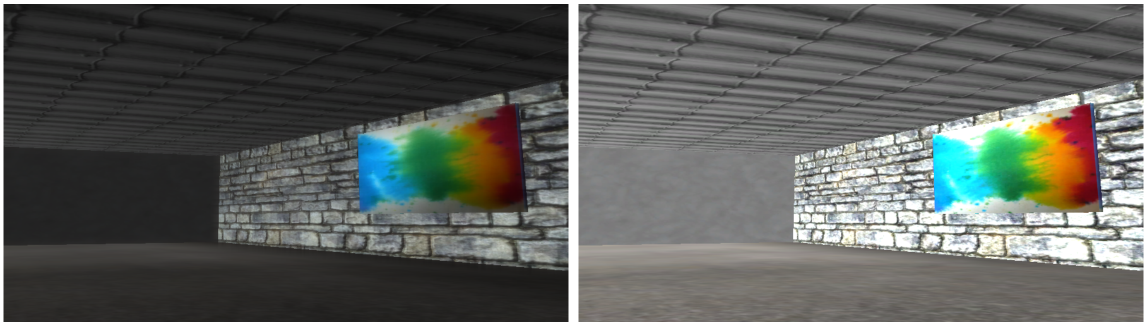

Figure 7.2: The source image (left) and the result of the cat light sensitivity transformation (right). The cat visual system has a lower minimum light detection threshold, meaning it can function in situations where light is more scarce. The intensities of the lights in the scene are adjusted by taking the logarithm of the light intensities perceived by a human and mulitplying it by a constant that represents cats' superior light sensitivity relative to humans.

7.2.2 Bee Light Sensitivity Transformation

Bees have a much higher minimum light detection threshold than humans, and as a result they cannot see well in the dark [10]. The VDS franework represents this difference by using a linear transformation that reduces the intensities of all the lights in the scene. A reduction of 30% seems appropriate to convey the impression that bees have more difficulty than humans under low light conditions. Figure 7.3 shows the result of the bee light sensitivity transformation.

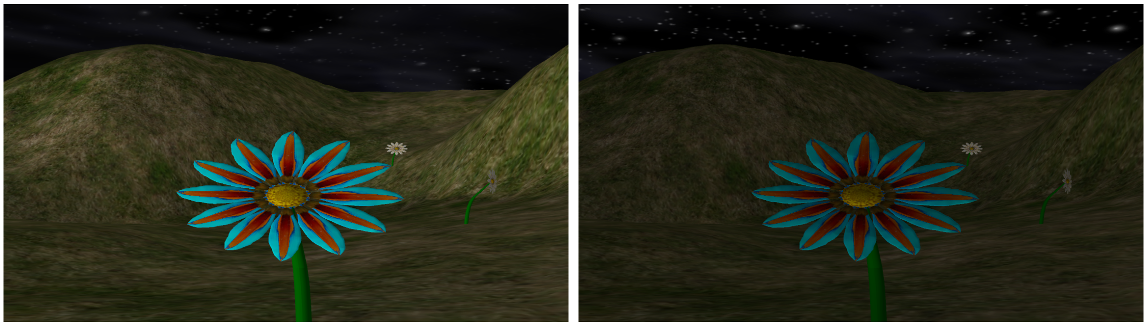

Figure 7.3: The source image (left) and the result of the bee light sensitivity transformation (right). Bees have inferior night vision compared to humans, and the simulation represents this by reducing the intensities of all the lights in the scene by 30%.

7.2.3 Pit Viper Light Sensitivity Transformation

Pit viper vision is optimized for light sensitivity and motion detection at the expense of visual acuity. Consequently, pit vipers have a lower minimum light detection threshold. The VDS framework can represent this using a logarithm relation that increases low intensities of light, but plateaus the intensities at higher levels. Figure 7.4 shows the result of this transformation.

Figure 7.4: The source image (left) and the result of the pit viper light sensitivity transformation (right). Pit vipers have superior night vision to humans, and the simulation represents this using a logarithmic function that increases light at low levels, but plateaus as light grows more abundant.

Chapter 8

Motion Sensitivity

Most visual systems are quite sensitive to movement, as detecting motion can be critical for revealing nearby predators or prey. The human visual system is more optimized for foraging, which requires the ability to focus and detect fine details, rather than motion tracking [43]. Most animals do not need such fine-grained vision. They sacrifice precision in exchange for characteristics like field of view or motion detection. Some animals, such as cats, have a *visual streak* of concentrated photoreceptors, rather than a dense foveal region like humans, and that makes them more sensitive to motion along the horizontal plane [55, 65, 68]. The *visual streak* is an asset when attempting to catch mice and other prey.

The VDS framework assumes that motion in the computer-generated environment is the baseline detected by a human observer. The goal of this transformation is to represent visual systems that have greater motion sensitivity—to convey in a visual manner that these systems are more adept at detecting movement than the human visual system. A classic approach for emphasizing fast motion is to use speed lines—lines that follow moving objects and indicate their direction of motion [50], but this is generally a two-dimensional approach that is not as common in situations where realism is desired. Motion blur has been employed in 3D rendering to make motion appear more realistic, as it takes into account the integration that a camera does when observing a moving target. However, motion blur can make it more difficult to ascertain an object's current position, as it becomes more indistinct. Instead, the simulation draws inspiration from traditional animation, where motion is sometimes emphasized by using 'ghosting'—a technique where moving objects leave behind a trail of repeated images, often more successively blurred [38].

The VDS framework achieves this effect by designating moving objects to leave a motion trail: a progressively more transparent record of their previous locations during the last few renders. This avoids blurring the moving object's current position, and makes it easier to detect moving objects and trace their recent path of movement. This is accomplished using a separate render pass to create a mask for all the moving objects that are currently visible. The data from this mask is added to a texture that contains the combined motion masks over the last several renders, with an alpha channel that deprecates over time at a rate determined by the system's motion sensitivity relative to human motion sensitivity. The motion mask texture is alpha-blended with the final render of the scene to add motion trails to moving objects that fade over time. The VDS framework requires that a system's relative motion sensitivity be input in the 0 - 1.0 range, and this value is multiplied with the alpha channel of the motion mask at each render. The higher the motion sensitivity value, the longer motion trails will persist. A motion sensitivity of 0 is equivalent to baseline human motion sensitivity, where moving objects leave no trails, while a motion sensitivity of 1.0 leads to long motion trails that fade very slowly.

8.1 Evaluating the Motion Sensitivity Transformation

8.1.1 Semantics

Representing motion sensitivity in a meaningful manner is challenging, since it is another characteristic that is visually ambiguous from a human perspective. As such, using transparency to represent this characteristic helps maintain consistency with the semantics of the hyperspectral transformation described in Chapter 5. The VDS framework incorporates transparency as part of a ghosting effect, where moving objects are designated to leave behind a trail of repeated images from the last several renders, each more successively transparent. The technique of using repeated images as a form of illustrating motion has been used extensively in traditional animation [38], so it already holds semantic value for representing movement.

8.1.2 Independence

The transparency in the motion sensitivity transformation has the potential to interfere with other instances where the same sort of technique is employed, such as with the hyperspectral transformation described in Chapter 5, and the visual acuity transformation described in Chapter 6. However, the blur used for the motion sensitivity transformation is linked to a moving object, and so is easier isolate from effects that are present regardless of the dynamics of the entities in the scene.

8.1.3 Efficiency

The motion sensitivity transformation requires a separate render pass to generate a mask for all the elements in the scene that are moving. The masking step is necessary to disambiguate moving objects from the movements of the virtual camera. Next, the new motion mask is combined with the motion masks from the last several renders, while deprecating the alpha values in accordance to the passage of time. A final render pass is used to alpha-blend the combined motion mask with the render of the scene. Even without further optimization, the motion tracking transformation is capable of running in real time on commodity graphics hardware.

8.2 Motion Sensitivity Example

8.2.1 Pit Viper Motion Sensitivity

Figure 8.1 shows a comparison between human vision and a pit viper's motion sensitivity. The small bouncing ball is almost impossible to detect from the human perspective, but is easier to observe from the perspective of the pit viper, due to the motion trail that it leaves behind. The dynamics of the motion sensitivity transformation make it even more noticeable when running in real time, but the difference is still perceptible in the static images shown in Figure 8.1.

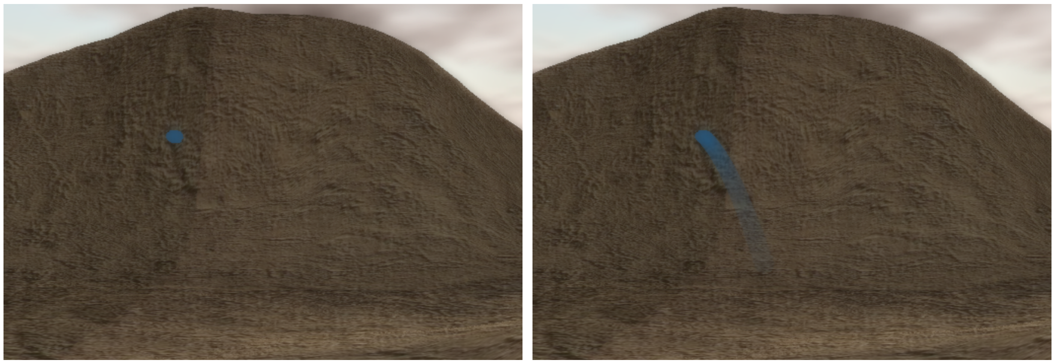

Figure 8.1: The motion tracking transformation is used to represent visual systems that are more sensitive to detecting motion. The human visual system has trouble detecting the small bouncing blue ball (left), whereas the pit viper simulation uses a ghosting effect that shows a motion trail behind the moving object, making it easier to notice [right].

Chapter 9

Field of View

Field of view (FoV) refers to the scope of the perception produced by a visual system—the extent of the world that is seen. A wider field of view allows a species to see more of their surroundings at any given moment. This characteristic is primarily determined by eye placement and construction. Species like humans with frontally placed eyes have a smaller field of view, but a larger region of binocular vision where the view offered by the eyes overlap. Species with laterally placed eyes have a much wider field of view, but their vision is mostly monocular, leading to more limited depth perception.

FoV is difficult to represent in a computer-generated environment, where the view of the scene is determined by the characteristics of the virtual camera displaying the scene. It is challenging to show differences in depth perception in what is already a monocular viewing model. Differences in the extent of the field of view are easier to represent, and are achieved in the VDS framework by changing the viewing angle of the virtual camera being used to observe the scene.

In keeping with the *difference-based* methodology, the VDS framework aims to display differences in field of view relative to a known baseline. A viewing angle of 45 degrees is standard for many virtual cameras, and this is used as the baseline for typical human vision. The framework represents the FoV of other visual systems by modifying the camera viewing angle based on how their FoV relates to human FoV. Animals with a wider field of view are thus able to see a more expansive view of the world, although as the viewing angle grows very wide it begins to cause distortion effects.

Equation 9.1 shows how the viewing angle of the virtual camera θ is calculated. It is based on σ, the field of view of the visual system being simulated, specified in degrees. ϕ is the typical human FoV specified in degrees, while d is a constant used to exaggerate the differences between human FoV and σ, the FoV of the system being simulated. Setting d to 0.75 seems to succeed in emphasizing the differences in FoV. The denominator ensures that a camera viewing angle of 45 degrees is used to represent ϕ, the standard human FoV. The viewing angle of the virtual camera, θ, will generally fall in the range of [27 - 90].

$$\theta = (\sigma + (\sigma - \phi) * d)/(\phi/45) \tag{9.1}$$

The VDS framework further emphasizes the differences in FoV by altering the size of the viewport used to display the simulation. The VDS framework generates the default human view so it includes black vertical margins at the edges of the viewport, and the width of these margins is altered depending on the relation between the human FoV and the FoV of the visual system being simulated. This captures the semantics that some species perceive a ‘widescreen’ view of the world when compared to humans. Equation 9.2 shows how w, the width of the margins in the human view, is adjusted to produce w', a new width for the margins that takes into account the FoV of the system being simulated. σ is the field of view of the system being simulated expressed in degrees, while ϕ is the standard human FoV. The framework uses a w of 1/8 of the screen’s width when representing normal human vision. w' will range from [0 - $w * 2$].

$$w' = w * 2 - w * \sigma/\phi \tag{9.2}$$

9.1 Evaluating the Field of View Transformation

9.1.1 Semantics

In principle, changing the viewing angle of the virtual camera used to view the scene offers meaningful semantics to represent the idea that different visual system perceive different extents of the world. However, this alone would remain one of the more subtle effects in the simulation. In addition, the VDS framework represents the idea that species with wider fields of view are getting a ‘widescreen’ view of the world, as opposed to those with more narrow fields of view. The framework accomplishes this by changing the width of the viewport based on whether the animal’s field of view is wider or narrower than the typical human field of view.

9.1.2 Independence

The FoV transformation changes the perspective matrix that is used for rendering the 3D scene in a manner to suggest differences in peripheral vision and field of view. While this will affect the way the scene is rendered, it should not have much impact on the results of the other transformations. It will not specifically affect the colors of objects in the scene, nor their brightness. It will change the way the scene attenuates over distance, but the objects in the scene will retain their characteristics.

9.1.3 Efficiency

The FoV transformation leaves little footprint on the efficiency of the simulation. Changing the width of the viewport and the viewing angle of the camera are both constant time operations, and do not need to be repeated except when the user decides to switch to simulating a different visual system. Consequently, their effect on efficiency is negligible.

9.2 Field of View Example

9.2.1 Cat Field of View

Cats are thought to have a wider field of view than humans, estimated at around 200 degrees. This gives cats more expansive peripheral vision [33]. This characteristic is simulated by changing the projection matrix of the camera when rendering the virtual scene from the cat's perspective, as shown in Figure 9.1. The viewport is also widened to represent that the cat is getting a 'widescreen' view of the environment.

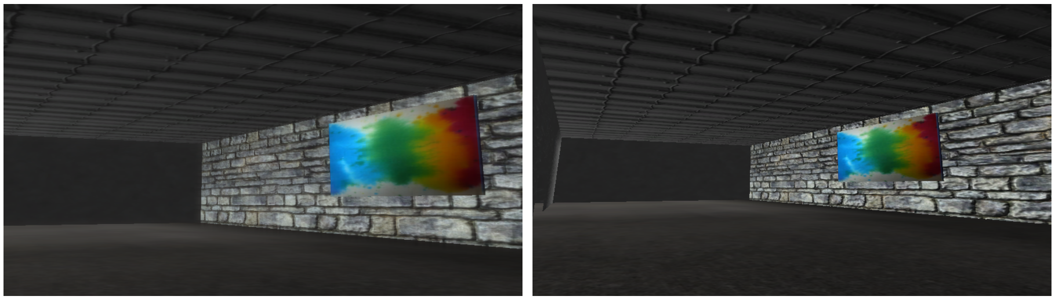

Figure 9.1: The source image (left) and the result of the cat field of view transformation (right). The cat field of view transformation renders the scene with a wider viewing angle, and shrinks the margins of the viewport to give a wider view of the scene.

Chapter 10

Eye Placement and Construction

Some species have compound eyes, where their visual field is produced by many photoreceptor and lens units, each aimed in a slightly different direction. In some instances, these views are integrated as part of their visual processing pipeline [83], but the VDS framework aims to show these views before integration to make the contrast in structure more clear. Once again, this is a case where the simulation prioritizes the illustration of differences in a meaningful manner over physiological accuracy.

Portraying thousands of separate views on a single display would be prohibitive both for the graphics processor, and for the sake of human comprehension. Instead, the VDS framework presents a small subset of the viewports, amounting to several dozen simultaneous views. Rather than rendering each in a separate viewport, which leads to considerable overhead, the framework achieves the compound vision transformation using indirect textures [41]. Indirect textures store texture coordinates at each pixel, which are used to index into a second texture. With this approach, each draw cycle needs only a pair of texture lookups to represent compound eye construction.

As a pre-processing step, an indirect texture is generated based on the structure of the compound eye being simulated. Appropriate texture coordinates are computed to provide the necessary overlap between each photoreceptor. The indirect texture is then used to index into a full-screen render of the scene. The result conveys the appearance of overlapping viewports at a fraction of the cost of rendering dozens of overlapping views of the same scene.

The user can specify the structure of the compound visual system in a configuration file, defining the shapes of the photoreceptors (hexagons and rectangles are currently supported), the amount of the screen that they occupy, the spacing between them, and their overlap in terms of texture coordinates. The VDS framework

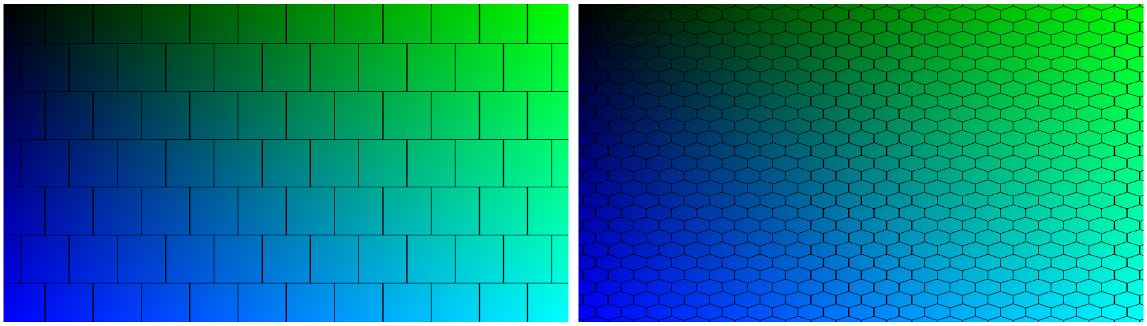

Figure 10.1: Indirect textures generated to represent different sorts of compound eye construction, including rectangular viewports (left) and hexagonal viewports (right). The indirect textures contain texture coordinates at each pixel location that can be used to index into a full screen render of the scene. The texture coordinates are visualized using a colormap where the green values represent the x texture coordinate and the blue values represent the y texture coordinate at a given pixel.

then automatically generates the indirect texture needed to produce the effect based on those specifications. Figure 10.1 shows examples of indirect textures generated to represent rectangular and hexagonal viewports with specified dimensions and overlap.

The indirect texturing approach used by the VDS framework is an approximation of a multiple viewport configuration, and some precision is sacrificed by using a constant overlap between photoreceptors in texture space rather than world space, but the results seem compelling enough to convey the desired impression to the observer. Figure 10.2 compares indirect texturing with a multiple viewport approach, where each rectangle is rendered as a separate viewport. The comparison was conducted on a 2.40GHz Intel quad-core CPU with 3.25 GB of RAM and a GeForce8600 graphics card. This example includes 97 viewing rectangles of 75x75 pixels each. The unmodified input runs at 1300 frames per second (FPS), while the multiple viewport implementation runs at 4 FPS—insufficient to maintain real time interaction. The indirect texturing approach runs at 202 FPS. The fourth image in each row shows a difference image comparing the output between multiple viewports and indirect texturing, and shows that there is little visual difference between them.

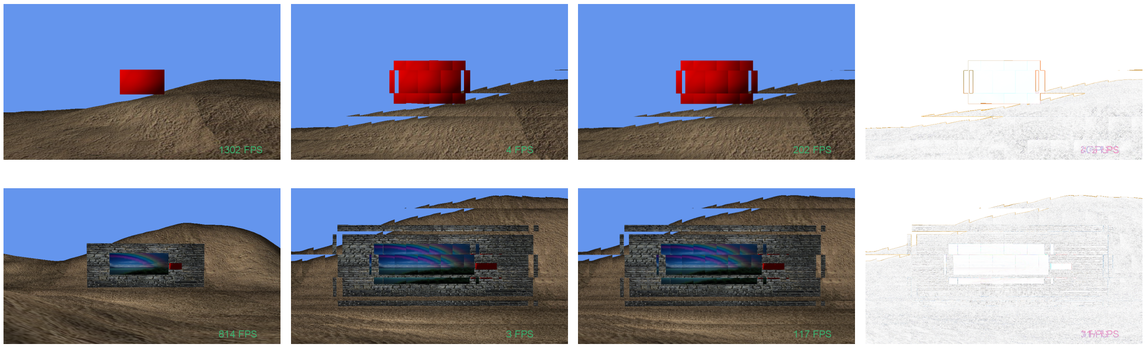

(e) Input (f) Multiple Viewports (g) Indirect Texturing (h) Difference Image

Figure 10.2: A comparison of multiple viewports against indirect texturing on a simple input scene using 97 viewing rectangles of 75x75 pixels each. The unmodified input [left] runs at 1300 frames per second (FPS), while the approach where each rectangle is rendered as a separate viewport runs at a non-interactive rate of 4 FPS. The indirect texturing approach runs at 202 FPS. A difference image (right) shows that there are only minor visual differences between the output produced by the two approaches.

10.1 Evaluating the Compound View Transformation

10.1.1 Semantics

Representing compound eye construction requires an illustration of the fact that the visual system in question is integrating views from many receptor units aimed in different directions. The VDS framework makes this explicit by rendering a representative number of separate views, even though in reality some animals with compound eyes integrate the signals from their many photoreceptors into a more blurred version of a single view.

10.1.2 Independence

The compound vision transformation primarily affects the view of the scene, and so it will have an impact on the field of view transformation described in Chapter 9. Compound eye construction is one reason that some visual systems have a substantially different field of view, so it stands to reason that it should be linked semantically to the field of view transformation. Although some animals with compound eyes integrate their views into a single blurred view, keeping the compound views separate helps this transformation remain distinct from the visual acuity transformation described in Chapter 6.

10.1.3 Efficiency

Thanks to the indirect texturing approach, the VDS framework is able to illustrate compound eye construction in real time with a non-trivial number of ‘viewports’. Producing the indirect texture as a pre-process takes a few seconds, but at runtime the visual system is rendered efficiently in real time. After the scene is rendered as normal, the compound eye transformation requires only a pair of texture lookups to effectively represent compound vision.

10.2 Compound Vision Example

10.2.1 Bee Compound Vision

Bees have compound eyes that are made up of thousands of roughly hexagonal photoreceptor units called *ommatidia*, each aimed in a slightly different direction [82]. The views from these photoreceptors are integrated at a later point in their visual processing pipeline. The VDS framework shows some of the separate views produced by the compound eyes before integration in order to more effectively illustrate this part of the bee’s visual system. Figure 10.3 shows the source image followed by the same scene when viewed through the simulation of bee compound eyes. The resulting effect conveys the impression of looking out simultaneously through dozens of different viewports that face in slightly different directions.

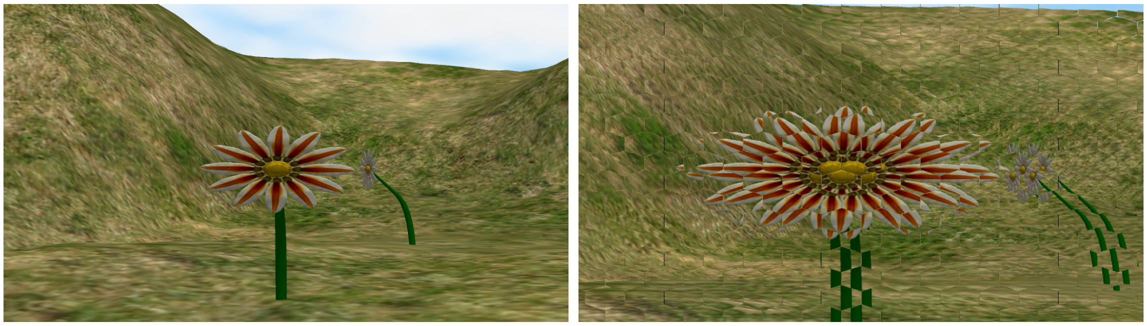

Figure 10.3: The source image (left) and the result of the bee compound eye representation (right). Bee compound eyes include thousands of *ommatidia*, each aimed in a slightly different direction. The framework represents this by generating an indirect texture to represent the overlapping hexagonal viewports as part of a pre-processing step, and then using the indirect texture at run-time to index into a full screen render of the scene.

Chapter 11

Case Study 1 - Cat Vision

The VDS framework produces illustrative simulations of visual systems based on how they relate to human vision. The input required to drive the transformations that make up the VDS framework are most easily harvested from within a computer-generated environment, and this means that the simulations are well suited to being incorporated within educational computer games. Computer games are increasingly seen as a valid form of experiential learning, both because they can be made to embody sound learning and cognitive principles, and because they can motivate and stimulate learners to take a more *active* role in the learning process [22]. By designing a game around the material present in a simulation, the process of learning what the simulation has to teach can potentially be made more interesting, more effective, and more *engaging*.

This section describes the first of two case studies that evaluate the educational impact of the simulations produced by the VDS framework. Each case study discusses a game that was designed to emphasize the elements of a particular visual simulation. In the process, several pedagogical principles important in the game-based learning literature are examined. Both games are implemented in C# using Microsoft's XNA game development library. They require Pixel Shader 3.0 support, and can be expected to run in real time on any machine with a graphics card at least equivalent to an NVIDIA GeForce 8600.

The first case study is based on a simulation of the cat visual system. Cats were chosen as the subject for this case study both because they are very popular as pets in North America, and their visual system is relatively similar to human vision, so it can serve as a starting point for representing a small set of visual differences and incorporating them into the mechanics of an educational game. Figure 3.1 shows the cat simulation produced by the VDS framework, and Appendix A documents the parameters used to create it. Section 11.1 describes *Catalyst*, the game that was developed to emphasize the differences between human and cat vision, while Section 11.2 discusses how this game was evaluated in terms of engagement, interest, and educational impact.

11.1 Catalyst Game Design

Catalyst is a first-person puzzle game intended to emphasize the differences between the human and cat visual systems. During the game, the player must switch between controlling a human and a cat in order to make their way past a series of obstacles to reach the goal. While controlling the cat, the players sees the world through the cat simulation. The puzzles that must be overcome are specifically designed to require players to switch between human and cat vision, thus emphasizing the differences between them. The design of the game was based on the principle of situated cognition [18] that suggests the information being learned should be directly applicable to the situation in which it is introduced. The structure of the game is shown in Figure 11.1.

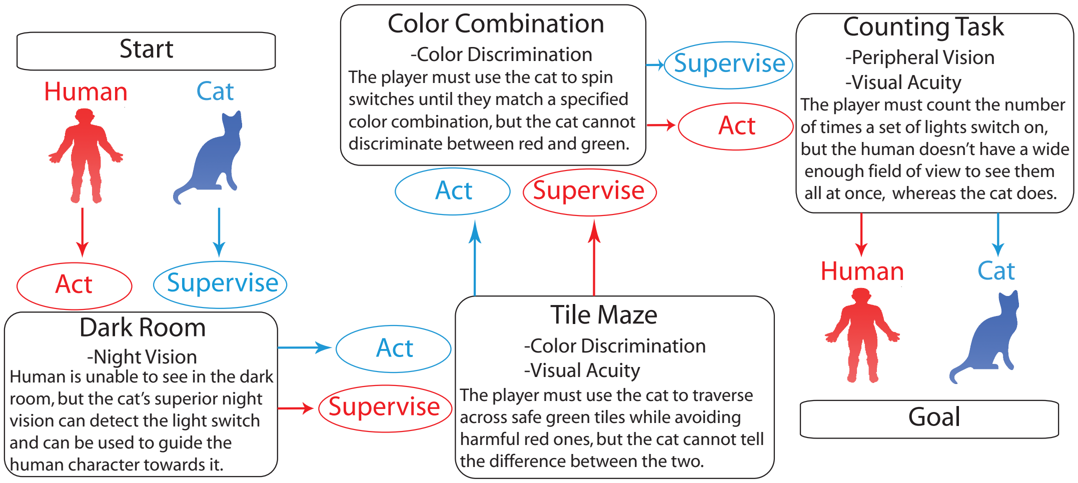

Figure 11.1: *Catalyst* is a first-person puzzle game that is designed to emphasize the differences between human and cat vision. The player must switch between controlling a human and a cat in order to overcome a series of obstacles. This diagram shows the obstacles in the game, along with the visual characteristics they are intended to emphasize. The obstacles in *Catalyst* have been set up so as to encourage the player to switch repeatedly between the two characters in order to use the eyes of one to supervise and guide the actions of the other. More description of the puzzles can be found in Section 11.1.

The first puzzle the player faces is based on low light vision. The player is initially looking through the eyes of the human, and can see very little of the surrounding environment. By switching to the cat character, the player is able to take advantage of the cat's superior night vision to locate the light switch in the corner of the room. The cat cannot reach the light switch, so the player must switch back and forth between the two characters to guide the human to the switch so it can be activated.

The next two puzzles that confront the player are based on color discrimination. In the first case, the player must guide the cat across a field of safe green tiles while avoiding harmful red ones. The cat cannot distinguish between the red and green colors, so the player must switch periodically back to the human character to supervise the cat's movement across the field of tiles. In the second case, the player must input a color combination to unlock a door.

The final puzzle is intended to emphasize the differences in field of view between the two visual systems. The player is asked to monitor a set of lights, and count the number of times they blink on and off. Due to the spacing between the lights, the human character cannot see all of the lights at once, making the counting task more difficult. However, the cat has a wider field of view, and can thus keep an eye on all the lights at once.

Each of the puzzles described above is designed to emphasize one of the major visual differences between cats and humans. The mechanics encourage the player to switch back and forth between the two visual systems, repeatedly contrasting them against each other when exposed to the same stimuli. These obstacles will thus encourage the player to build a mental mapping between the two visual systems.

11.2 Experimental Design

It is challenging to evaluate casual educational games where the term of engagement is limited to a few minutes of play, and the parcel of knowledge being conveyed is consequently small. The *Catalyst* evaluation was designed to compare the different methods that could be used to convey the four differences between the human and cat visual systems displayed by the simulation. The classic approach to convey such information in interpretive settings like zoos or museums would be plain text, perhaps augmented by images or video clips. As such, the experiment compares playing *Catalyst* against watching a video of the game being played, or reading a text document containing facts about cat vision.

The two main research questions were:

- How does the change in interest compare between text, game, and video learning tasks that all cover the same material?

- How does the knowledge acquisition compare between a game and video playback of the same game? Does the interaction involved in a game lead to a more comprehensible and engaging learning experience?

18 subjects participated in the evaluation with University ethics approval. The subjects were gathered from the university community, and included undergraduate students, graduate students, and professors. It should be noted that this sample does not necessarily reflect the results one might obtain from the general population, but it can serve as a proof-of-concept that the simulations have the educational potential. Table 11.1 shows the experimental design. The 18 participants were divided randomly into two groups of nine so that two experimental conditions could be tested in a between-subjects design. Both groups began the experiment with a pre-test, where they responded to Likert scales intended to gauge their interest in cats and their visual systems, and their experience with computer games. All questions were posed in a random order to avoid context effects, and approximately half were reverse-coded to avoid acquiescence bias. Respondents were also asked to write down any facts they already knew about the cat visual system so their a priori knowledge could be separated from the knowledge they gleaned during the trial.

After the pre-test, members of the first group of participants played through the game, *Catalyst*, while the members of the second group watched a video playthrough of the game. In both cases, the learning task was followed by post-test questions. The interest questions from the pre-test were repeated so a delta could be extracted that represents the participants' change in interest during the trial. The post-test also included Likert items intended to assess the participants' level of engagement with the learning task. Two dimensions of O'Brien's model of engagement (endurability and involvement) were used as the instrument for this measurement [57]. Participants were also asked to write down any characteristics of the cat visual system that they observed during the learning task.

Both groups were asked to complete a second learning task as part of a within-subjects design where they would read a document that described the four factors of the cat visual system emphasized in the simulation and game. The post-test questions following this final task asked the participants to rate their engagement while reading the text, and to directly compare the two tasks they had completed in terms of educational value, engagement, and stimulation. In so doing, participants were able to give information on how they perceived the game and video tasks as compared to the conventional means of conveying information through plain text. This task was

Evaluation	**Game Condition**	**Video Condition**
Participants	9	9
Pre-Test	Likert items in Appendix C.1 regarding: • Computer experience • Interest in animal vision Open-ended knowledge question in Appendix C	Likert items in Appendix C.1 regarding: • Computer experience • Interest in animal vision Open-ended knowledge question in Appendix C
Task 1	Playing the *Catalyst* game	Watching a video of a *Catalyst* playthrough
Post-Test	Likert items in Appendix C.2 regarding: • Interest in animal vision • Engagement during Task 1 Open-ended knowledge question in Appendix C	Likert items in Appendix C.2 regarding: • Interest in animal vision • Engagement during Task 1 Open-ended knowledge question in Appendix C
Task 2	Reading instructive text shown in Appendix C.2	Reading instructive text shown in Appendix C.2
Post-Test 2	Likert items in Appendix C.3 regarding: • Engagement during Task 2 • Direct comparison of Task 1 and Task 2	Likert items in Appendix C.3 regarding: • Engagement during Task 2 • Direct comparison of Task 1 and Task 2

Table 11.1: The experimental design used to evaluate *Catalyst* and the cat illustrative simulation produced by the VDS framework. Participants were divided into two groups as part of a between-subjects design to compare those who played the game against those who watched the video. The game and video were also compared against reading informative text as part of a within-subjects design. Section 11.3 describes the results of this evaluation.

performed second in both experimental cases so as to focus on testing the knowledge acquired from the game and the video without interference from the text task. As a result, there may be ordering effects present in the data associated with this task. Appendix C contains more specific details on the instruments used for this evaluation.

11.3 Results and Discussion

Figure 11.2 shows how the participants' interest in cats and their visual systems changed after playing the game or watching the video. Participants recorded a significant increase in interest in both the game condition (M = 0.78, median = 0.75, s = 0.54) and the video condition (M = 0.53, median = 0.5, s = 0.63), according to Wilcoxon tests ($p < 0.05$). However, a Mann-Whitney test suggests no significant difference between the interest increases caused by the game and the video. Several respondents later commented that they had grown more intrigued about cat vision following the test, and one respondent mentioned that playing *Catalyst* had given him the urge to check online information sources like Wikipedia as soon as he finished the test so that he could learn more about the subject.

Figure 11.3 shows that the participants in the game condition felt more engaged while playing the game (M = 84.7222, median = 85, s = 11.689) than while reading the text document (M = 62.2225, median = 70, s = 14.1115). A Wilcoxon test suggests that this difference is statistically significant ($p < 0.05$). The participants in the video condition showed no significant difference in engagement between watching the video (M = 66.39, median = 60, s = 20.3527) and reading the text (M = 69.1675, median = 67.5, s = 8.6602). A Mann-Whitney test comparing the two conditions suggests that the game was considered significantly more engaging than the video ($U = 63, p < 0.05, r = 0.4683$).

The engagement results imply that the participants who completed the game task felt that they had become more involved in the task, while those watching the video felt more like passive participants and were not as actively engaged in their task. The participants in the video task may also have felt more disconnected because they had no control over the progression of their task, or the pace at which it was conducted, leading to confusion. Two respondents who watched the video mentioned that they had difficulty interpreting what was being shown to them. The missing connection that comes from interacting with the game and seeing the resulting feedback in the virtual environment made it difficult for them to distinguish between random events and those triggered by a specific action.

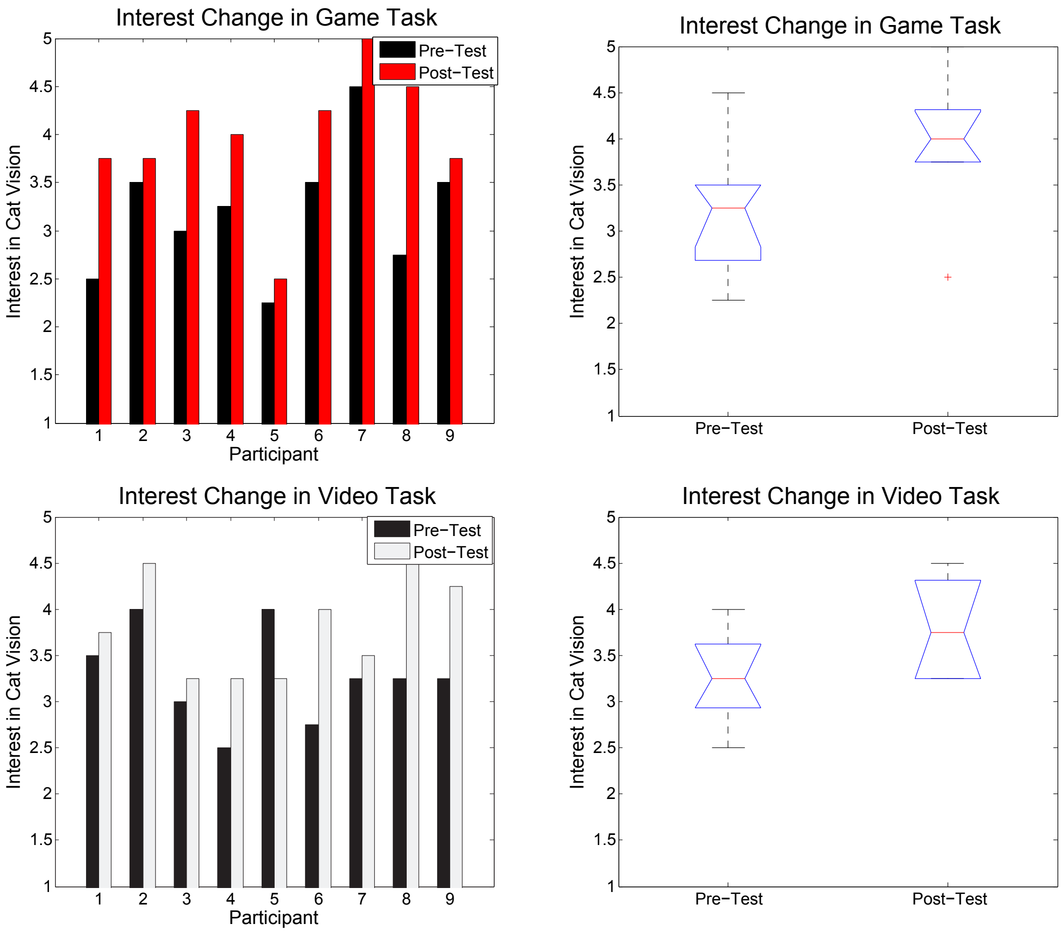

Figure 11.2: The interest change reported by participants in the game condition (top), and the video condition (bottom). Both tasks produced a statistically significant increase in interest among participants according to Wilcoxon tests ($p < 0.05$), although a Mann-Whitney test suggests no significant difference between the interest increases caused by the game and video conditions.

Figure 11.4 shows how the participants responded when asked to directly compare the game or video tasks against the text reading text. The game group was significantly more interested in the game over the text (M = 33.3375, median = 62.5, s = 47.18625) according to a binomial test ($p < 0.05$). However, a binomial test suggests that there was no significant preference for the video over the text (M = -36.1112, median = -12.5, s = 47.78). One possible interpretation of this data is that the participants that played the game felt they had accomplished something substantial, while those watching the video felt like passive participants without a tangible goal, leading to a lesser sense of accomplishment upon the completion of the task. However, it should be noted that there could be ordering effects present in this data, since in both conditions the text task was performed second.

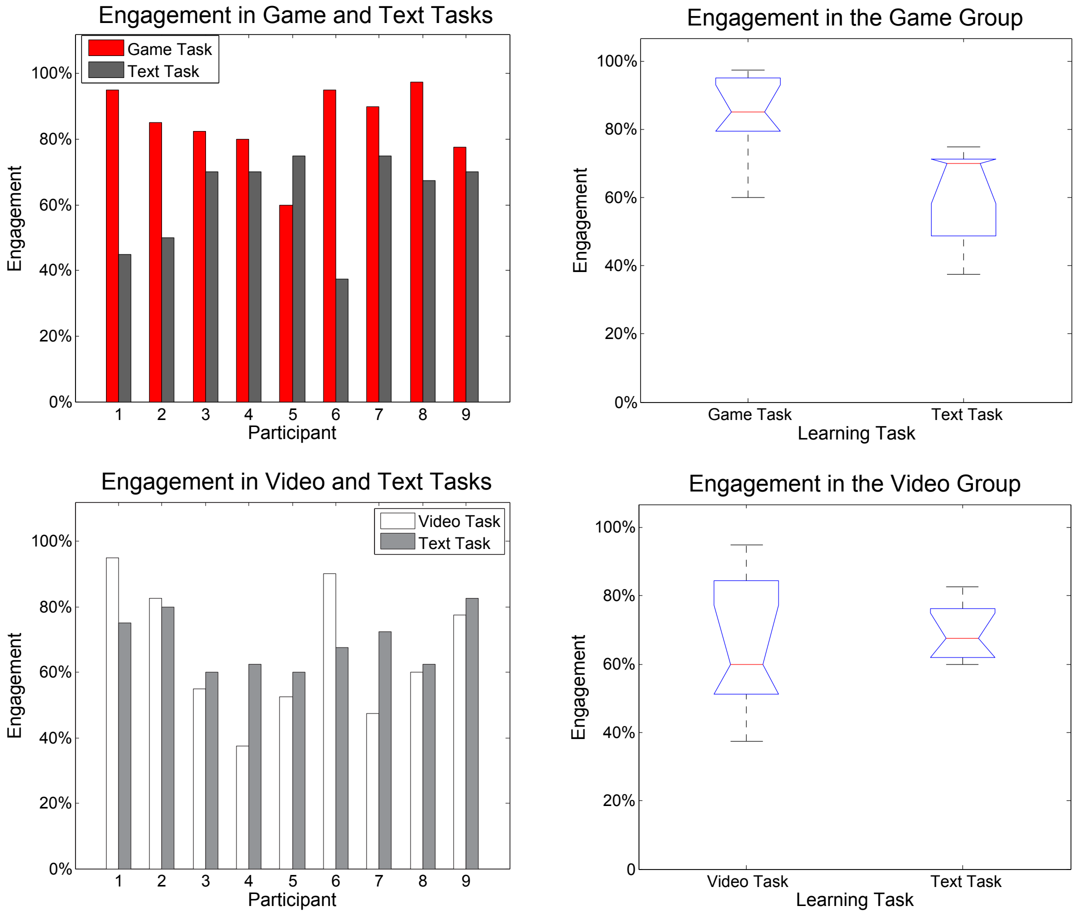

Figure 11.3: The engagement of the participants in the two different experimental conditions. The participants in the first group felt significantly more engaged while playing the game than while reading the text (top), while those in the second group did not feel significantly more engaged while watching the video than while reading the text (bottom). A Mann-Whitney test indicates the group playing the game was significantly more engaged than the group watching the video ($p < 0.05$).

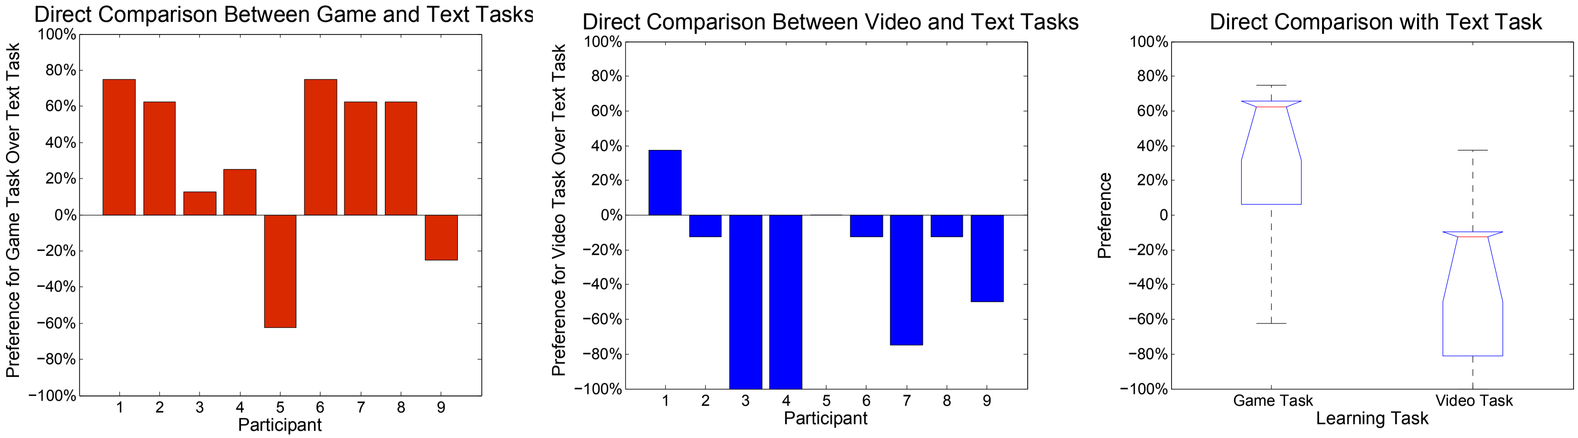

Figure 11.4: The participants in the group that played the game showed a statistically significant preference for it as a learning tool over a traditional text document according to a binomial test ($p < 0.05$), whereas the second group showed no significant preference between the video and traditional text as a learning tool. Positive results on this figure indicate a preference for the game or video over the text, while negative results indicate a preference for the text.

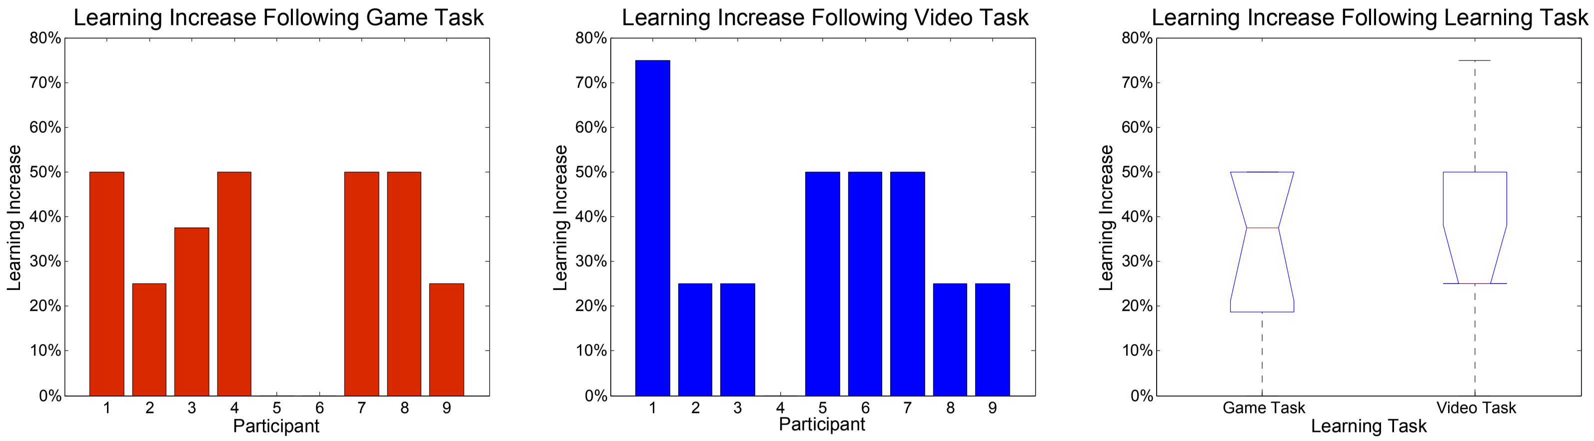

Figure 11.5: Both conditions in the *Catalyst* experiment produced statistically significant increases in knowledge according to Wilcoxon tests ($p < 0.05$), although there is no significant difference between the game and video group. Table 11.2 shows the learning increases broken up by visual characteristic rather than by participant.

Figure 11.5 shows the learning increases of the participants who played the game (M = 31.94, median = 37.5, s = 20.83) and those who watched the video (M = 36.11, median = 25, s = 22.05). The data in this figure shows the percentage of the four major differences between humans and cats that the participants noticed during the learning task. Wilcoxon tests indicate that both groups achieved statistically significant learning increases ($p < 0.05$), although a Mann-Whitney tests suggests no significant difference between the increases recorded by the game and video groups. This is not inconsistent with previous research, as it has been suggested that one of the advantages of game-based learning is long-term retention of information, even in cases where no immediate benefit is perceived [58, 61, 64]. However, it should be noted that the freeform instrument used to assess learning in this experiment (shown in Appendix C.2) made it difficult to determine the depth of learning achieved. Participants were often vague in reporting the characteristics that they had noticed, and this motivated the use of a more formalized instrument to assess learning in the second case study described in Chapter 12.

Table 11.2 shows the learning increases reported in the *Catalyst* experiment broken up by visual characteristic. The color and night vision aspects of the cat visual system were widely acknowledged by both groups, but the cat's poor visual acuity and expanded field of view went mostly unobserved by the participants. The failure of the participants to notice the poor visual acuity of the cat visual system speaks to the importance of situated cognition within educational games. The acuity transformation was not directly relevant to any of the puzzles that the players needed to overcome to complete the game, and likely went unnoticed as a result. Game players tend to be very focused on achieving their goal, to the exclusion of elements that are irrelevant to their task. This is consistent with the findings of perception researchers with regard to inattentional blindness [71]. Development of future educational games should recognize that information may not be retained if it is superfluous to the tasks in the game.

The participants who played the game showed a significant increase in knowledge regarding the cat's expanded field of view according to a Wilcoxon test ($p < 0.05$), unlike those who watched the video. The field of view transformation is one of the more subtle features of the cat simulation, and the light counting obstacle that is intended to demonstrate the cat's expanded field of view is more complicated than the other obstacles in the game. The participants watching the video were not forced to decipher and overcome this obstacle in order to complete their learning task, and as a result they may not have fully understood its significance. The participants playing

	Video			**Game**		
Characteristic	**Pre-**	**Post-**	**Delta**	**Pre-**	**Post-**	**Delta**
Color Vision	**33.89%**	**88.89%**	**55%**	**33.33%**	**100%**	**66.67%**
Night Vision	**50%**	**88.89%**	**38.89%**	**66.67%**	**88.89%**	**22.22%**
Field of View	**0%**	**33.33%**	**33.33%**	**0%**	**38.89%**	**38.89%**
Visual Acuity	**0%**	**22.22%**	**22.22%**	**0%**	**0%**	**0%**
Mean	**20.97%**	**58.33%**	**36.11%**	**25%**	**56.95%**	**31.95%**

Table 11.2: The learning increases reported by participants in the *Catalyst* experiment broken up by visual characteristic. A Mann-Whitney test suggests no significant difference between the learning increases produced by the two different conditions. The color and night vision characteristics were universally acknowledged, but few noticed the differences in visual acuity or field of view.

the game, on the other hand, could only complete their task by overcoming the light counting obstacle, and so they were forced to come to some understanding of it.

According to the evaluation that was conducted, *Catalyst* was successful both in stimulating further interest in the subject matter, and in conveying knowledge to the players. In particular, the educational material that was presented using a situated cognition approach—where it was directly relevant to the player's progression through the game—was more likely to be retained, whereas the material that had no impact on the player's advancement through the game was less likely to be noticed.

Chapter 12

Case Study 2 - Bee Vision

The *Catalyst* case study serves as a compelling proof of concept. Based on its success, the VDS framework was extended to include more diverse visual differences. Bees were chosen as the next species to simulate, as they exhibit several more challenging visual characteristics, such as hyperspectral sensitivity and compound eye construction. Figure 3.3 shows the bee simulation produced by the VDS framework, and Appendix A documents the parameters used to create it. Section 12.1 describes the three part evaluation that was conducted to compare the educational benefits of the bee simulation as a stand-alone learning tool versus the benefits of incorporating the simulation into a game scenario.

12.1 Experimental Design

The *Catalyst* experiment had demonstrated the importance of performing an evaluation before beginning to design the game that would incorporate the bee simulation, in order to test if the simulation as a stand-alone has value as an educational tool. As such, the evaluation of the bee simulation was divided into three stages, each one involving a user study with different participants based around the same set of Likert items. The participants were all from the university community, and it should be noted that this sample does not necessarily reflect the results one might obtain from the general population, but it can serve as a proof-of-concept that the simulations have the educational potential. The common testing instrument provides a basis for comparing the three stages of the evaluation against one another using a between-subjects design. There was a period of appromixately two months between each stage of the experiment, and different subjects participated in each one.

- Stage 1 - Simulation Experiment: the user is asked to interact with the simulation by exploring a computer-generated environment in an unstructured fashion;
- Stage 2 - Performance task: the user is asked to interact with the simulation while participating in a simple performance task within a computer-generated environment;
- Stage 3 - Bee Prepared: the user plays an educational computer game that is designed to incorporate the bee simulation and emphasize the information present within it.

The first challenge was to come up with a testing instrument that could be used as a common basis for evaluating all three stages of the bee vision experiment. Each experiment would include some evaluation metrics specific to it (such as performance task data for stage 2, and gameplay data for stage 3), but these would all be augmentations of the common testing instrument. In the end, twelve Likert items were devised that focus on the five major differences between human and bee vision. The participants were asked to respond to this set of Likert items both before and after each learning task, so that a delta could be extracted measuring how their knowledge had changed due to the task. All the Likert items measured the participants' agreement with correct and incorrect statements about bee vision. Participants' knowledge was determined by how strongly they agreed with correct statements and disagreed with incorrect ones. Appendix D includes more specifics on the common instrument for this experiment. Table 12.1 shows the experimental design used for this case study.

12.2 Simulation Experiment

Before embedding the simulation in a game scenario, it was important to test the simulation as a stand-alone learning tool—without any of the seductive details that can come from burying educational content in the guise of a game. Ten respondents from the university community participated in the simulation experiment. First, the participants were given a pre-test with the standard testing instrument shown in Appendix D.1. Next, the participants were given fifteen minutes to explore a computer-generated environment, with the ability to switch between a human and bee perspective at the push of a button. The first couple minutes of the experiment included a short automated tutorial explaining how to interact with the system, but the remainder of the task was unstructured and unsupervised.

Evaluation	**Simulation**	**Performance**	**Game**
Participants	10	14	20
Pre-Test	Likert items in Appendix D.1 regarding: • Differences between human and bee vision	Likert items in Appendix D.1 regarding: • Differences between human and bee vision	Likert items in Appendix D.1 and D.3 regarding: • Differences between human and bee vision • Interest in bee vision • Computer familiarity
Task 1	Interacting with the bee simulation for 15 minutes	Performance task under one of two conditions.	Playing *Bee Prepared* for 15 minutes
Post-Test	Likert items in Appendix D.1 regarding: • Differences between human and bee vision	Likert items in Appendix D.1 regarding: • Differences between human and bee vision	Likert items in Appendix D.1 and D.3 regarding: • Differences between human and bee vision • Interest in bee vision • Enjoyment
Task 2	None	Performance task under the condition that was not tested in Task 1.	None
Post-Test 2	None	Likert items in Appendix D.1 and D.2 regarding: • Differences between human and bee vision • Comparison of Task 1 and Task 2	None

Table 12.1: The experimental design used to evaluate the bee illustrative simulation produced by the VDS framework. The experiment was conducted in three stages so that a between-subjects design could be used to compare the bee simulation as a stand-alone learning tool against the educational impact of a game designed to incorporate it.

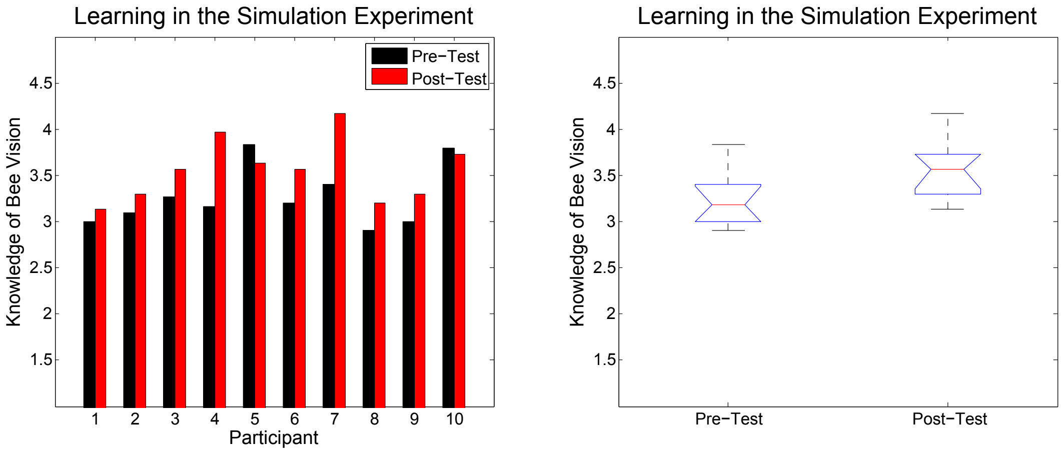

Figure 12.1: The mean results per participant on the pre- and post-tests for the simulation experiment. The results are shown on a five point scale, where higher numbers indicate greater knowledge of bee vision. The participants interacted with the simulation in an unstructured fashion for fifteen minutes, and this led to a statistically significant increase in knowledge about the bee visual system according to a Wilcoxon test ($p < 0.05$).

Figure 12.1 shows the mean learning changes recorded between the pre- and post-tests for the simulation experiment (M = 0.29, median = 0.30, s = 0.32). A Wilcoxon test indicates that the interaction with the bee simulation led to a statistically significant increase in knowledge about bee vision ($T = 4.5, p < 0.05$). The results of this stage of the evaluation suggest that the simulation as a stand-alone can be a compelling learning tool. The next challenge was to embed it in a task-oriented scenario without compromising its educational impact.

12.3 Performance Task

The simulation experiment suggests that the bee simulation has some merit as a stand-alone educational instrument, but letting the user interact with it in their own way does not guarantee a consistent experience. It may not motivate the user to explore all the ideas presented by the simulation. The second stage of the evaluation was designed to go one step further to investigate two questions:

- Will giving the user a task to perform have a noticeable effect on the learning outcomes of the simulation?
- Is a glow effect more suitable for representing the ultraviolet light that bees can detect than a traditional false coloring approach?

Two options were considered as a means for the VDS framework to represent hyperspectral sensitivity. This experiment was partially motivated as a means of comparing the effectiveness of the glow effect described in Chapter 5 relative to a more traditional false coloring approach. Which manner of representing bees' hyperspectral sensitivity would be more meaningful to unprimed observers in context with all the other effects?

A performance task was devised where users were shown flowers of different colors within the bee simulation, and asked to distinguish between them and indicate what colors they would correspond to in the human color space. The flowers were specified to reflect ultraviolet light, and in the first condition that ultraviolet light was represented with false coloring using a purple colormap, whereas in the second condition the ultraviolet contribution was represented using the glow effect described in Chapter 5.

The performance task specifically involved showing the user a set of three flowers and asking them how many of the flowers would appear a certain color from a human perspective. Five different colors were used, with two different ultraviolet contributions: weak, and strong. The respondents were shown flowers in random sets of three, and eight sets were shown for each of the five colors present in the experiment. In total, each participant inspected 120 flowers during a run of the experiment. Figure 12.2 shows some screen shots from the performance task of the same scenes where in one case the glow effect is used to represent ultraviolet light, and in the other false coloring is employed instead.

Fourteen new respondents participated in the performance experiment using a within-subjects experimental design where each respondent engaged in the performance task under both the glow and false coloring conditions. The participants were randomly divided into two groups. The first group began with the glow condition before moving on to the false coloring condition, while the second group worked in the opposite order. In addition to gathering performance data on the number of correct identifications and responding to the standard testing instrument developed for all three stages of the evaluation, respondents were also quizzed in terms of their preferences between the two methods for representing the hyperspectral data.

Figure 12.3 shows the average number of color misidentifications that the participants made under the glow condition (M = 21.67, median = 22.5, s = 8.42) and the false coloring condition (M = 21.43, median = 18.75, s = 6.23). A Wilcoxon test shows no statistically significant difference between the two conditions in terms of the average number of color misidentifications.

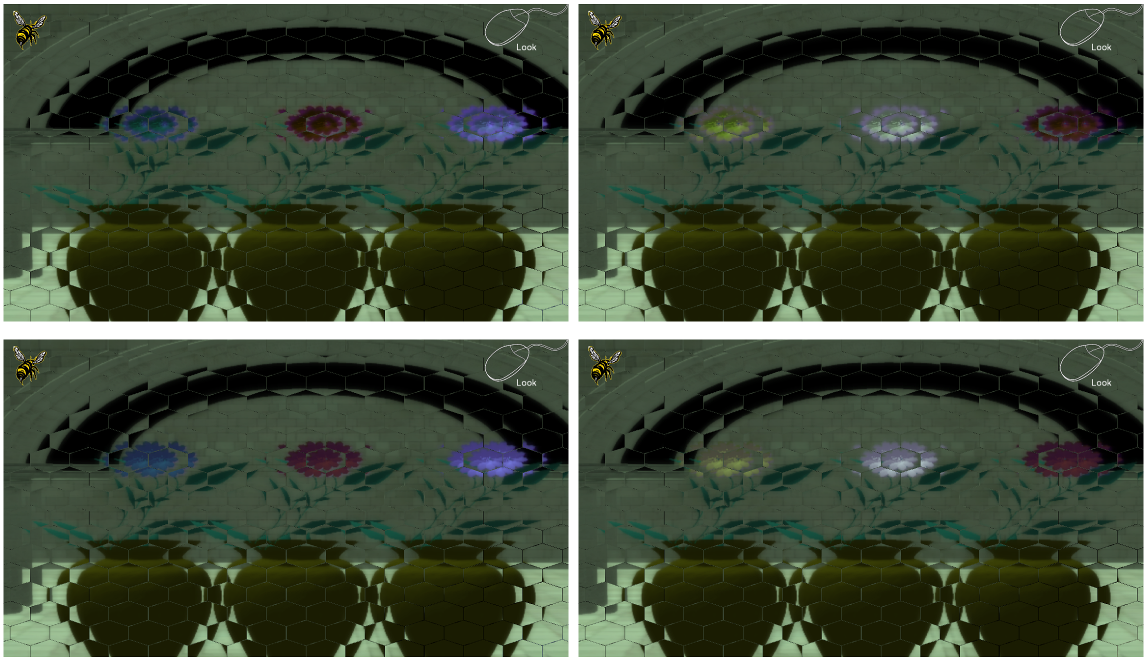

Figure 12.2: Screenshots from the performance task under the glow condition [top] and the false coloring condition [bottom]. While respondents did not show a statistically different learning outcome between the two conditions overall, the performance task suggests that as contributions of ultraviolet light grow stronger, respondents have more difficulty with color identification under the false coloring condition than under the glow effect.

One area of interest is to consider the impact of the two different UV intensities represented in the experiment. Figures 12.4 and 12.5 show the false negatives for each respondent, broken up by UV contribution. Wilcoxon tests indicate there is no significant difference between the number of false negatives in the glow condition between weak UV contributions (M = 17.74, median = 17.5, s = 15.19) and strong UV contributions (M = 26.16, median = 28.4, s = 16.84). However, a Wilcoxon test suggests there are significantly more false negatives in the false coloring condition when there is a strong UV contribution (M = 34.63, median = 35, s = 19.82) than when there is a weak UV contribution (M = 12.76, median = 15, s = 8.62) ($T = 2, p < 0.05$). This suggests that the false coloring is more likely to compromise the discriminatory power of the other colors in the scene as the UV contribution increases and the intensity of the false coloring dominates the colors present in the scene. This is problematic, as the colors are critical for conveying an important aspect of the visual system being simulated. The glow effect, on the other hand, seems to be perceived more independently from the colors in the scene. Higher UV contributions are represented by a brighter glow, but this does not have a significant impact on a

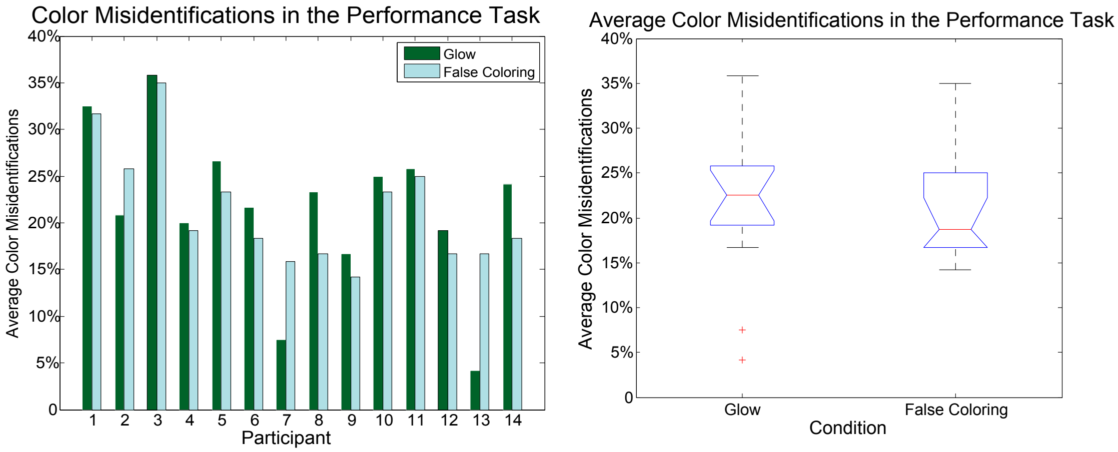

Figure 12.3: Average number of color misidentifications for each respondent under the two different conditions across the five different flower colors and the two different UV contributions. A Wilcoxon test suggests no significant difference between the average number of color misidentifications in the glow and false coloring conditions. There is also no significant ordering effect between the condition that respondents were subjected to first versus the one they participated in second.

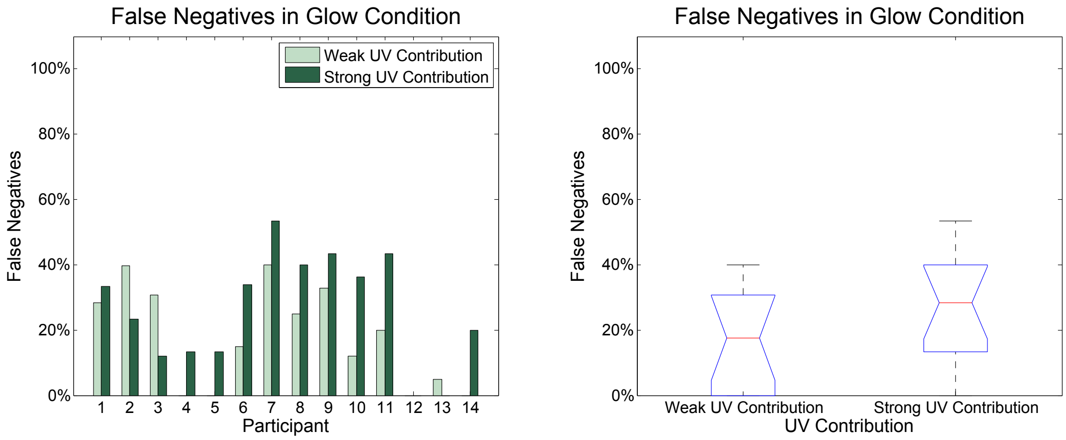

Figure 12.4: The number of false negatives reported by respondents during the performance task under the glow condition. A Wilcoxon test shows there is no significant difference under the glow condition between the number of false negatives with weak and strong UV contributions, suggesting that the UV had little impact on the respondents' ability to judge colors in the performance task under the glow condition.

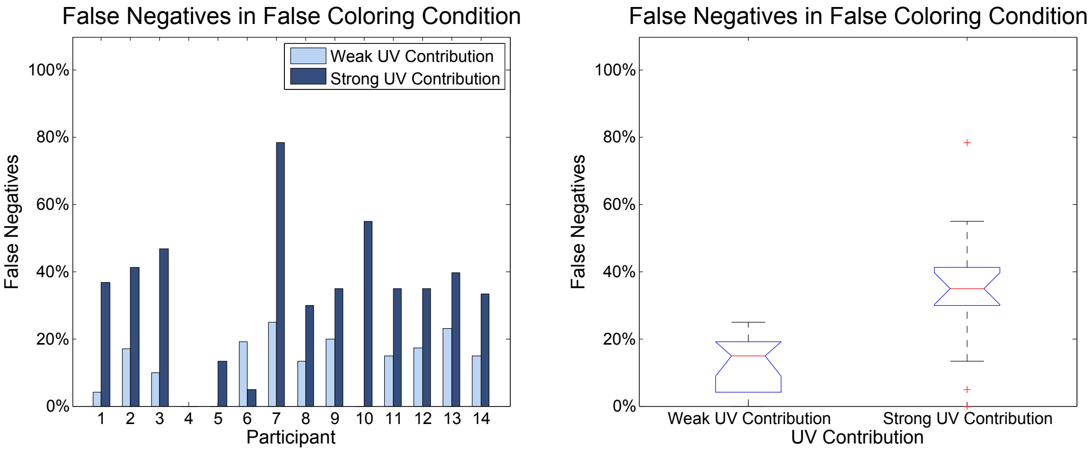

Figure 12.5: The number of false negatives reported by respondents during the performance task under the false coloring condition. A Wilcoxon test indicates a statistically significant difference between the false negatives identified when they exhibited a weak UV contribution and when they exhibited a strong contribution ($p < 0.05$). This suggests that more potent false coloring contributions make it more difficult for respondents to assess the underlying colors in the scene, compromising the impact of the simulation's color transformation.

respondent's ability to identify the underlying colors of the flowers in the scene.

Respondents were also quizzed in terms of their preference between the two methods for representing the hypersectral data, as shown in Figure 12.6. Participants were asked to directly compare the glow and false coloring effects in terms of preference and aesthetics. Their responses showed an insignificant preference (M = 17.86, median = 12.5, s = 45.92) for the glow effect, according to a binomial test. Respondents did not seem to associate the glow effect any more strongly with ultraviolet light, but they may have found it more appealing. All else being equal, it makes sense to select a representation that is perhaps more appealing.

Figure 12.7 shows the results of the common test instrument that indicate a statistically significant increase in knowledge regarding bees and their vision under both the glow condition (M = 0.42, median = 0.45, s = 0.30) and the false coloring condition (M = 0.49, median = 0.49, s = 0.30), according to Wilcoxon tests ($p < 0.05$). However, a Wilcoxon test suggests there is no statistically significant difference between the learning increase in the glow and false coloring conditions. Furthermore, a Mann-Whitney test indicates there is no significant difference between the learning increase produced by the performance task and the one seen in the simulation task, shown in Figure 12.1.

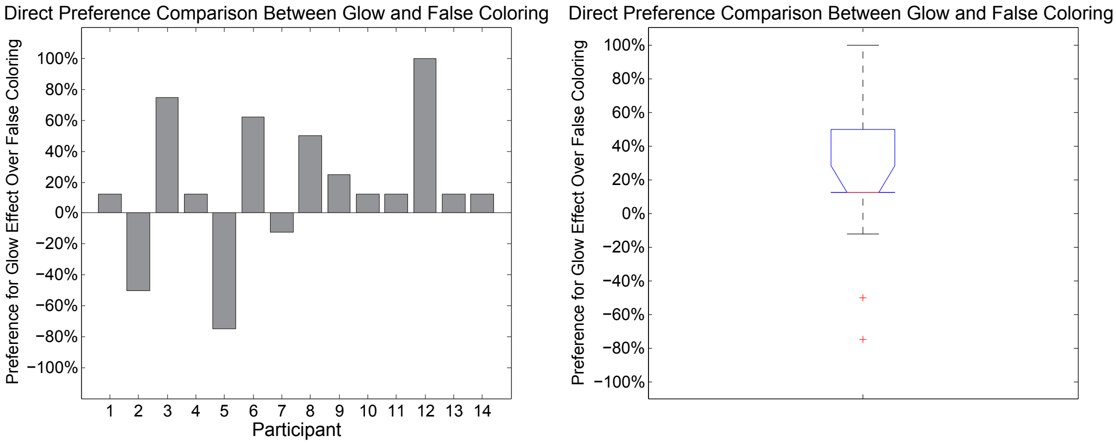

Figure 12.6: Participants in the performance experiment were presented with a series of Likert statements directly comparing the glow effect and false coloring in terms of aesthetics and personal preference. Data points on this graph greater than zero indicate a preference for the glow effect, while negative values signify a preference for the false coloring effect.

These results do not suggest that the performance task encouraged greater learning, but they do imply that it did not impede the bee simulation's ability to convey the educational content to the user. At the least, the experiment demonstrated that the performance task did no harm with regard to the learning outcomes. This encourages one to believe that it may be possible to further engage and motivate users by giving them tasks to complete, without diluting the impact of the simulation. Ideally, this property could be exploited to tap into the intrinsic appeal and motivation of a game while at the same time preserving the simulation's educational impact.

The last phase of the evaluation considered integrating the simulation into an educational game. The aim of the third phase was to see if embedding the educational content into the core mechanics of the game can not only avoid detracting from the learning outcomes, but can actually *enhance* the learning impact of the simulation.

12.4 Bee Prepared

The next sections discuss *Bee Prepared*, the game designed to incorporate and leverage the bee simulation. It is a first-person time-trial game where the player switches between controlling a human and a bee. The challenges in the game are carefully crafted based on cognitive principles so as to promote the educational material present in the bee simulation. The next section discusses these principles specifically, and how

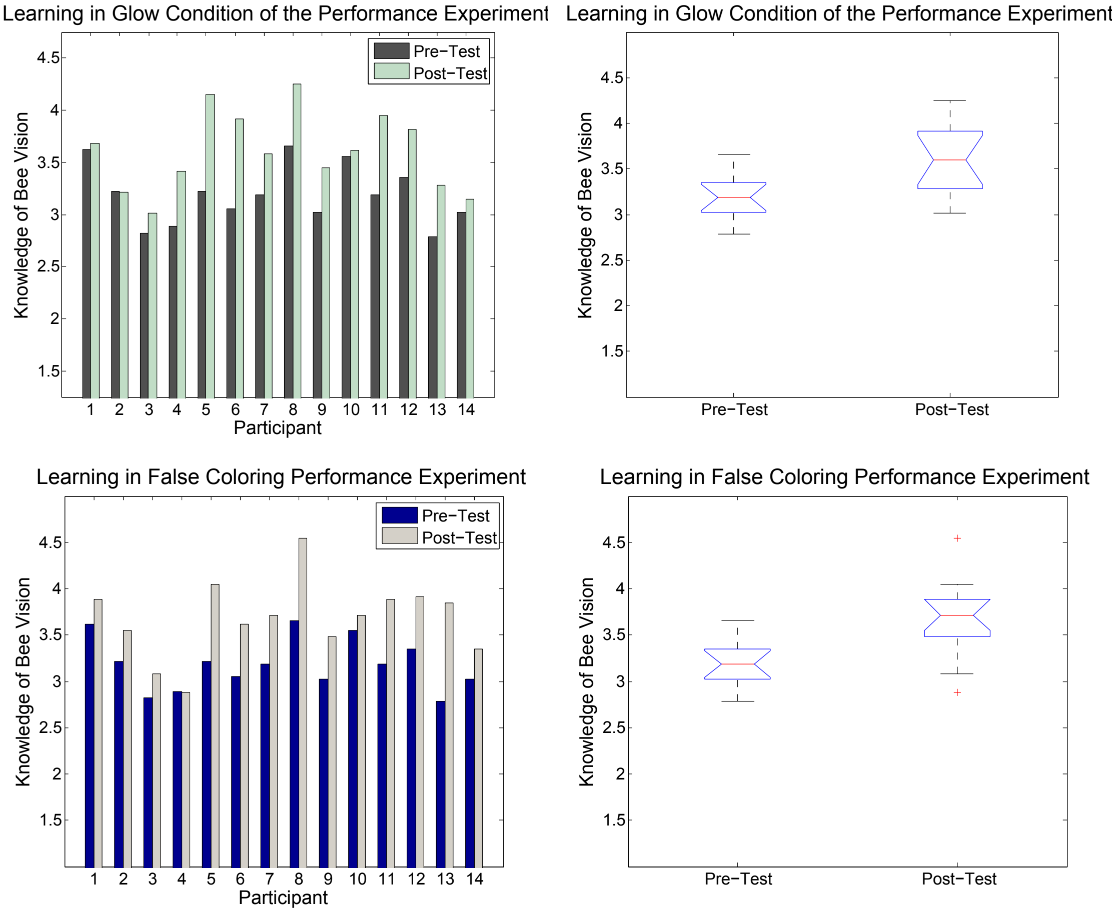

Figure 12.7: The mean results per participant on the pre- and post-tests for the performance experiment under the glow condition (top) and the false coloring condition (bottom). The results are shown on a five point scale, where higher numbers indicate greater knowledge of bee vision. Both conditions led to statistically significant increases in knowledge according to Wilcoxon tests ($p < 0.05$), but a Wilcoxon test shows no significant difference between the learning increases produced by the two different conditions.

they informed the game design decisions.

12.4.1 Design Factors

The first case study demonstrated the importance of situated cognition as a guiding principle for educational game design. This suggests that the material being taught should, in some sense, be immediately applicable to the tasks that the player is facing. In other words, the player should need to take the educational content into consideration in order to progress more effectively in the game, and to master the game mechanics. Gee referred to a similar idea as one of the learning principles present in games, although he described it as the *explicit information on-demand and just-in-time principle* [22]. The second guiding principle that was used in designing *Bee Prepared* was the idea of *repeated contrast* that suggests that differences should be emphasized through repetition, and that this repetition should be an important part of the gameplay, as suggested by Gee's *practice principle* [22].

Based on the principles described above, four main design goals were identified for *Bee Prepared*:

- Reinforce all the visual differences present in the bee simulation using as much situated cognition as possible. This requires designing game mechanics in such a way that they emphasize the differences between human and bee vision, and coming to a better understanding of these differences allows the player to achieve a comparative advantage.
- Give the player the ability to switch between human and bee view, and encourage them to do so repeatedly so they can more effectively build a mental mapping between the two visual systems.
- Try to integrate the learning material such that players can get a concentrated sample without necessarily having to get to the 'end' of the game.
- Keep the game mechanics as simple as possible. This lowers the barrier for entry into the game, and leaves the player free to contemplate how they can tailor their gameplay strategy in order to achieve a comparative advantage.

12.4.2 Bee Prepared Game Mechanics

Bee Prepared is a time challenge where the player must collect flower seeds of different colors within a time limit. The human character can plant seeds that will grow into flowers of the corresponding color. The bee character can pollinate these flowers, giving the player more seeds of that color. In reality, it would take considerable time for the fertilized flowers to produce seeds after being pollinated, and bees have no reason to ‘collect’ the seeds produced, but such simplifications help keep the game mechanics from growing too complex. The player can switch between controlling the human and the bee characters, although doing so incurs a time penalty and includes a cool down—a period of time that must elapse before the player is allowed to switch characters again. Randomly colored flowers will grow periodically as time goes on, so the player will have to discriminate between flowers they have planted and are looking to collect versus those that have appeared in the wild. Once the objective number of seeds is reached, a more difficult composition of seeds is presented as the new objective, and must be collected within a more limited time frame.

Figure 12.8 shows a visual representation of the core game mechanics. The central mechanic of the game is cyclical, forcing the player to switch repeatedly between human and bee in order to progress towards the goal, and this repeated contrast exposes them to the mapping between the human and bee perspectives. This relatively simple set of mechanics offers opportunities to emphasize the important visual differences between bees and humans, and to provide players with a comparative advantage if they take these characteristics into account in their gameplay strategies.

Figures 12.9 and 12.10 show screen shots from the game that have been annotated to identify the heads-up display (HUD) elements, and to give a sense of the player experience. Figure 12.9 shows the human perspective of the world that the player experiences while planting seeds of various colors in the flower pots scattered throughout the game world. Figure 12.10 shows the bee perspective that the player experiences while approaching and pollinating flowers. Note that the HUD includes a cool-down meter that tracks how much time must elapse before the player is allowed to switch to the other character again. Figure 12.11 shows the same scene from both the human and bee perspectives, and gives a side-by-side comparison of how their views of the game world differ.

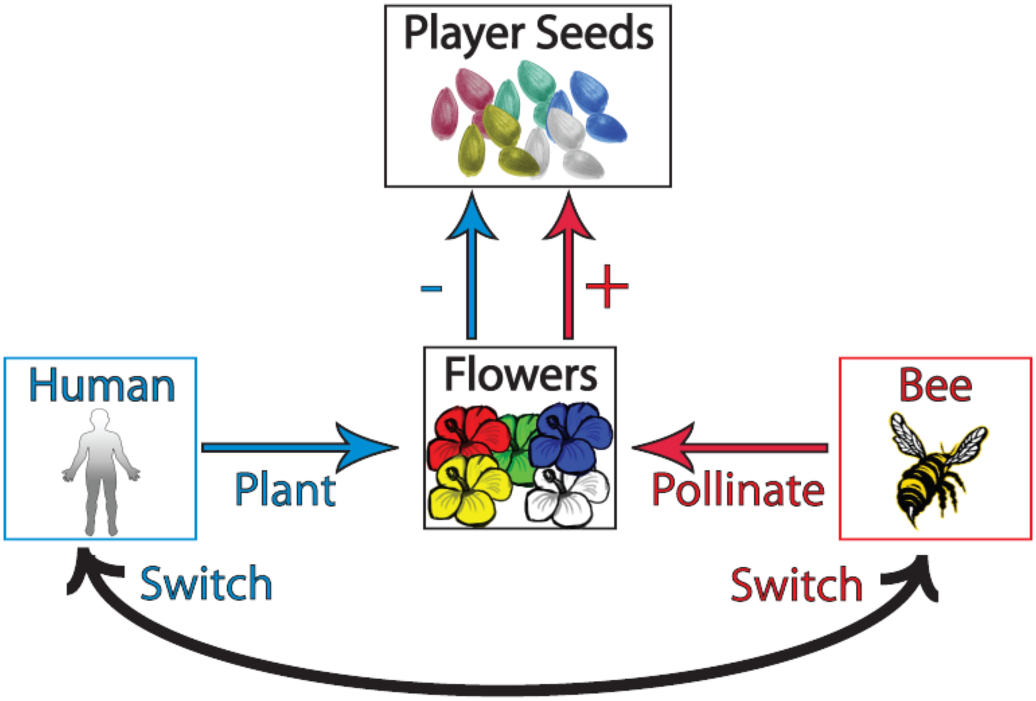

Figure 12.8: A simplified game design diagram for *Bee Prepared*. The player switches between controlling a human and bee character in order to collect a certain number of the five colors of flower seeds. The human character can plant seeds to create more flowers, while the bee character can pollinate planted flowers to earn more seeds. This cyclical mechanic encourages the player to continually switch between the two characters and perspectives in order to progress towards their objective.

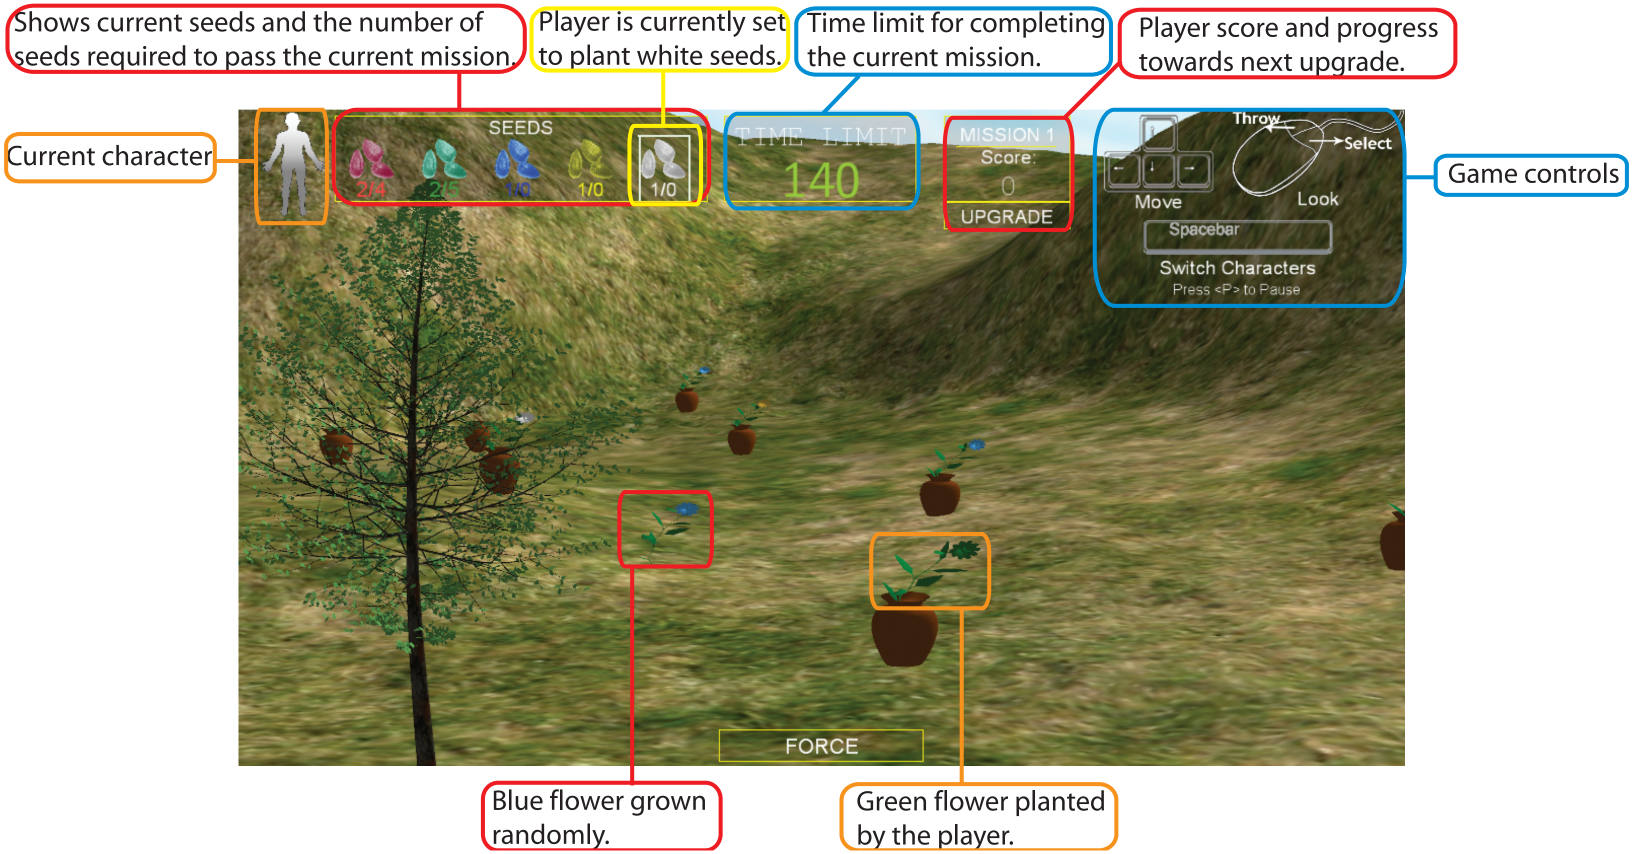

Figure 12.9: A screen shot from *Bee Prepared* showing the human character's view of the game world. The annotations describe the HUD (heads-up display) that the player sees, and are intended to convey some notion of the user experience while playing the game. While controlling the human, the player will attempt to plant flower seeds of the needed colors. The human's superior low light vision will make this a feasible task to perform at night, when the bee character is unable to see.

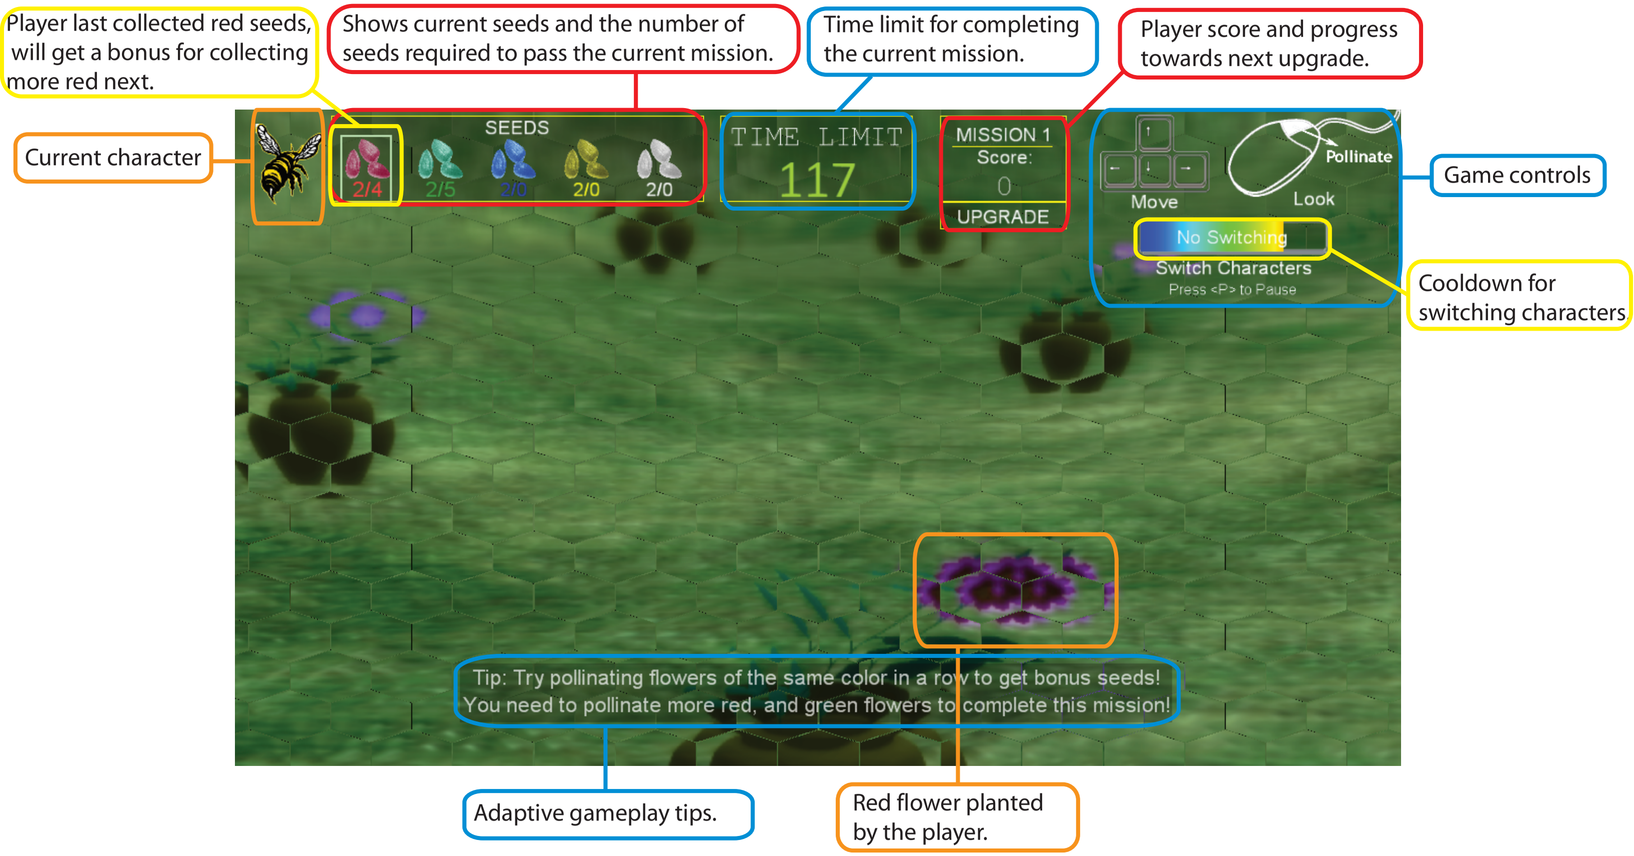

Tip: Try pollinating flowers of the same color in a row to get bonus seeds!
You need to pollinate more red, and green flowers to complete this mission!

Player last collected red seeds, will get a bonus for collecting more red next.

Current character

Shows current seeds and the number of seeds required to pass the current mission.

Time limit for completing the current mission.

Player score and progress towards next upgrade.

Game controls

Cooldown for switching characters

Adaptive gameplay tips.

Red flower planted by the player.

Figure 12.10: A screen shot from *Bee Prepared* showing the bee character's view of the game world. The annotations describe the HUD (heads-up display) that the player sees, and are intended to convey some notion of the user experience while playing the game. While controlling the bee, the player will attempt to approach and pollinate several flowers of the same color in a row, giving them more seeds of that type. This forces them to come to terms with the bees' color discrimination, visual acuity, and field of view. A cool-down factor on the ability to switch characters prevents the player from switching back to the human every time they need to determine the color of a flower.

Figure 12.11: Screen shots of *Bee Prepared*, the game developed to incorporate and enhance the bee simulation. The images show the same scene from the perspective of the human character (left) and the bee character (right). The game mechanics encourage the player to switch repeatedly between the two perspectives while playing, giving them an opportunity to build a mental mapping between the two visual systems.

The mechanics of *Bee Prepared* are designed to emphasize the five major differences between human and bee vision represented by the simulation. The following sections consider how the game mechanics encourage the player to take heed of each of these differences using situated cognition and repeated contrast.

12.4.3 Bee Color Vision

It is important to make it clear that the bee color vision system is more sensitive to lower wavelength colors, such as blues, purples, and ultraviolet light, and less sensitive to the higher wavelength range that includes reds and oranges. During the game, the player will need to control the human character to plant flowers, and then switch to the bee character in order to pollinate them. This means the player will be required to look at the same stimuli from the perspective of both visual systems in order to accomplish their task. This repeated contrast should help the player build a mental mapping between the 'human colors' they are used to and their correspondences in the 'bee color space'.

The player's objective is to collect a certain number of seeds of each color, so it will be important for them to be able to identify the colors of flowers from the bee's perspective before pollinating them. The color discrimination differences are further reinforced by including a mechanic whereby player receives more seeds for pollinating several flowers of the same color in a row (referred to in the game as a 'chain bonus'). To take maximal advantage of this bonus, players will need to be able to discriminate between the different colors of flowers from within the bee simulation. This situated cognition encourages the players to come to some understanding of the mapping between human and bee colors if they wish to do well at the game.

12.4.4 Ultraviolet Sensitivity

Bees are capable of perceiving ultraviolet wavelengths of light, and this sensitivity helps them locate and identify flowers. The player has a vested interest in successfully locating flowers quickly while controlling the bee character, and so they will reap benefits if they connect the glow effect to the flowers they are seeking. Given the other effects at play, the glow that represents hypserspectral data is one of the most visually noticeable. This is partially due to the dynamic nature of the glow, which attracts the player's attention. The player will thus receive a comparative advantage if they recognize that the glow is used to annotate flowers, but only from the bee's perspective.

12.4.5 Bee Eye Construction

Bee eyes are made up of thousands of *ommatdia*, each aimed in a slightly different direction. Bees' compound eyes give them a very different perspective of the world around them. The time pressure that the player is under means that they will be encouraged to navigate the world as quickly as possible, even while controlling the bee. Navigation will be markedly easier for the player if they come to some understanding of how the bee field of view differs from our own.

12.4.6 Bee Night Vision

Bee photoreceptors have a high minimum light detection threshold that makes it more difficult for them to see under low light conditions. *Bee Prepared* features a day/night cycle, causing the light levels to change drastically over time. The human character will be able to see better at night, and so the player will have a comparative advantage if they use the human to plant flowers during that period, and then switch to the bee to pollinate flowers during the day, reinforcing the notion that bees do not have very effective night vision.

12.4.7 Bee Visual Acuity

Bees have inferior visual acuity to humans, and this is actually one of the most difficult characteristics to convey, because it is challenging to design tasks that cater specifically to such a subtle feature of the visual system. In general, the time pressure the player is facing will encourage them to navigate as quickly as possible. Although the bee character is able to move more quickly, the inferior visual acuity—especially when looking over long distances—will make it difficult for the player to find and identify flowers that are far away.

12.4.8 Visual Upgrades

The mechanics described above are intended to emphasize the differences in five of the major visual characteristics between humans and bees. However, some of the effects are more subtle and difficult to convey in context with all the others, such as the visual acuity and night vision characteristics. To help emphasize these elements, and to give the player extra incentive to perform well, a simple upgrade system is integrated into the game, wherein the player is given an opportunity to choose a bonus between missions if they have achieved a high enough score. Figure 12.12 shows how the five upgrades are visually depicted in the game.

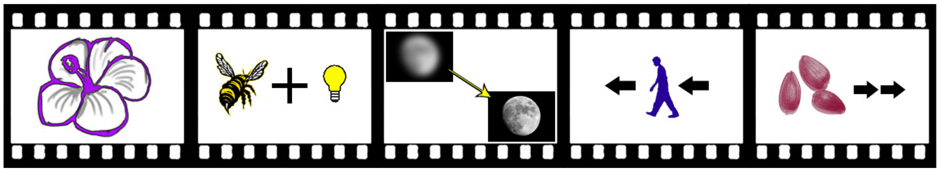

Figure 12.12: *Bee Prepared* gives players the ability to select more visual upgrades as they progress through the game. Each upgrade is intended to give the player an advantage in navigating through the world during upcoming missions. Three of the five available upgrades emphasize aspects of the bee visual system that are more difficult to convey without further explanation, such as ultraviolet sensitivity, ocelli light sensors, and nocturnal blur reduction. The player has an intrinsic interest in taking heed of these descriptions, as they will offer a tangible and immediate benefit in the missions to come.

Three of the five upgrades are linked to the visual differences being represented, including one that enhances ultraviolet sensitivity, one that reduces the blurriness that bees face during the night, and one that lowers the bee's minimum light detection threshold, making it easier for them to see at night. Two other upgrades are included for variety, including one that increases the player's speed and one that allows the human to throw and plant seeds from a greater distance. Each upgrade includes a paragraph of descriptive text, explaining how it fits into the game and into the bee's visual mechanisms. By linking this learning material into informing a choice that the player must make (in terms of which upgrade to purchase at any given opportunity), the player is then encouraged to take note of the learning material, as it should have a concrete impact on their ability to succeed in the game. This upgrade system also increases the game's re-playability, as the player will be able to experience a slightly different gameplay experience by choosing a different series of upgrades on their next playthrough.

12.5 Bee Prepared Evaluation

Bee Prepared was developed as a way of enhancing the educational impact of the bee simulation, which had already proven effective as a stand-alone in the first two phases of the evaluation process. The purpose of the third and final stage of the evaluation process was two-fold:

- Investigate whether embedding the bee simulation in a game would lead to more beneficial learning outcomes than presenting the simulation as a stand-alone;
- Determine if the game stimulates interest in the subject matter.

Twenty participants from the university community participated in this phase of the evaluation process. Participants were first asked to respond to Likert items regarding the differences between human and bee vision, establishing their pre-existing knowledge. This questionnaire included the standard test instrument developed and used throughout the first two phases of the evaluation process, in addition to items related to computer game experience, and interest in bees and their vision.

Following the pre-test, the participants were left to play *Bee Prepared* for fifteen minutes. They were given a short in-game tutorial describing how to navigate through the computer-generated environment and how to toggle between a ‘human’ and ‘bee’ view of the world around them. They also received instructions on how to play the game, including a short tutorial mission before proceeding to the timed levels. Gameplay data was logged while the participants played *Bee Prepared*, including statistics on how often they switched characters, achieved chain bonuses, the upgrades they selected, the level they reached during the experiment, and the time it took them to pass each level.

Following the learning experience, the participants were asked to respond to a post-test containing some of the same Likert items from the pre-test. Comparing their responses on the pre- and post-tests indicates how their knowledge of bee vision changed while completing the learning task. Finally, the study concluded by allowing respondents to indicate what they had observed and learned in their own words. Table 12.1 summarizes the experimental design used to evaluate *Bee Prepared*.

12.5.1 Results and Discussion

Figure 12.13 shows the mean learning increases observed by the 20 partipants in the game experiment. A Wilcoxon test suggests that the respondents recorded a statistically significant increase in knowledge after playing *Bee Prepared* ($T = 0, p < 0.05, r = -0.8771$). In considering Figure 12.13, one might observe that respondent 14 indicated no change in knowledge. This particular respondent spoke to the invigilator after the experiment and explained that although he felt the game had been attempting to convey information to him, he had nonetheless refused to change any of his answers because he did not know if he could ‘trust’ the information presented in the form of a game. It is worth noting that this kind of skepticism is still a force

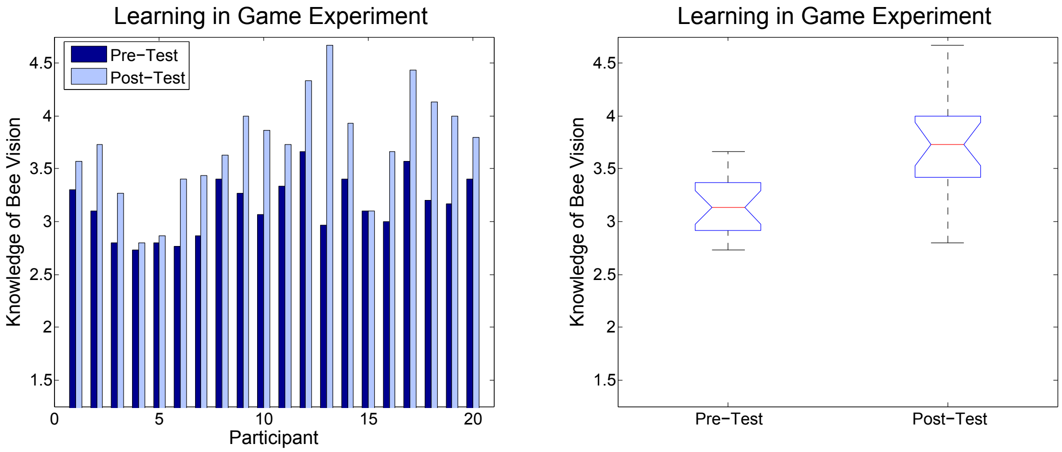

Figure 12.13: The mean results per participant on the pre- and post-tests for the game experiment. The results are shown on a five point scale, where higher numbers indicate greater knowledge of bee vision. A Wilcoxon test suggests a statistically significant difference between the results of the pre- and post-tests ($p < 0.05$), suggesting that the participants knew more about bee vision after playing the game. Overall, a Mann-Whitney test suggests that those who played the game had a significantly higher learning increase than those who interacted with the simulation as a stand-alone as part of the first two bee simulation experiments ($p < 0.05$).

to be reckoned with when employing game-based learning. His answers have been discounted from the statistical analysis because they did not appear to reflect upon this game in specific, but instead on games as a learning tool in general.

One of the objectives when conducting the evaluation was to determine if integrating the simulation into a game would be successful not just in being more engaging, but also in promoting the educational content. There is potential for the game elements of educational games to distract from the education, leading to a reduced learning outcome, and it is important to determine if the approach of embedding the learning material into the game's core mechanics could help avoid such a fate. A Mann-Whitney test suggests that the participants in the game experiment displayed a significantly higher learning increase (M = 0.57, median = 0.60, s = 0.38) than the learning increase achieved in the simulation and performance experiments, where the participants interacted with the simulation outisde a game scenario (M = 0.37, median = 0.33, s = 0.31) ($U = 317.5, p < 0.05, r = -0.4670$). The results suggest that the participants who played the game learned more than those who interacted with the simulation in a less structured fashion. In other words, the mechanics in the game seemed to be successful in emphasizing and *enhancing* the learning impact already present in the simulation.

	Simulation			Game		
Characteristic	**Pre-**	**Post-**	**Delta**	**Pre-**	**Post-**	**Delta**
Color Vision	2.97	3.07	0.10	2.95	3.13	0.18
Hyperspectral	3.19	3.33	0.14	3.05	3.67	**0.62**
Compound View	3.67	4	0.33	3.48	3.75	0.27
Visual Acuity	3.17	3.78	**0.61**	3.1	4.03	**0.93**
Light Sensitivity	3.1	3.75	**0.65**	3.15	4	**0.85**
Mean	3.22	3.59	**0.37**	3.15	3.72	**0.57**

Table 12.2: Participants in the experiment were asked to respond to the Likert items shown in Appendix D.1 that inquired about the five major differences between human and bee vision represented by the bee simulation. This table shows the mean responses broken up by visual characteristic on the pre- and post-tests. The responses are given on a five-point scale, where higher values indicate that the respondents demonstrated greater knowledge of that visual characteristic. Cells in bold indicate significant differences according to Wilcoxon tests ($p < 0.05$). Overall, a Mann-Whitney test suggests that those who played the game had a significantly higher learning increase than those who interacted with the simulation as a stand-alone ($p < 0.05$), as described in Figure 12.13. Figures 12.14 and 12.15 show box plots of learning increases broken up by visual characteristic.

Table 12.2 shows the results of the game experiment and compares them against the results from interacting with the bee simulation outside a game scenario. The simulation heading includes the results from both the simulation experiment described in Section 12.2 and the performance experiment under the glow condition described in Section 12.3. The data inTable 12.2 is broken down in terms of the five visual characteristics that are represented by the simulation. The table includes the mean result for each characteristic on a five point scale, where higher responses indicate greater knowledge of the characteristic. Figures 12.14 and 12.15 show box plots of learning increases broken up by visual characteristic.

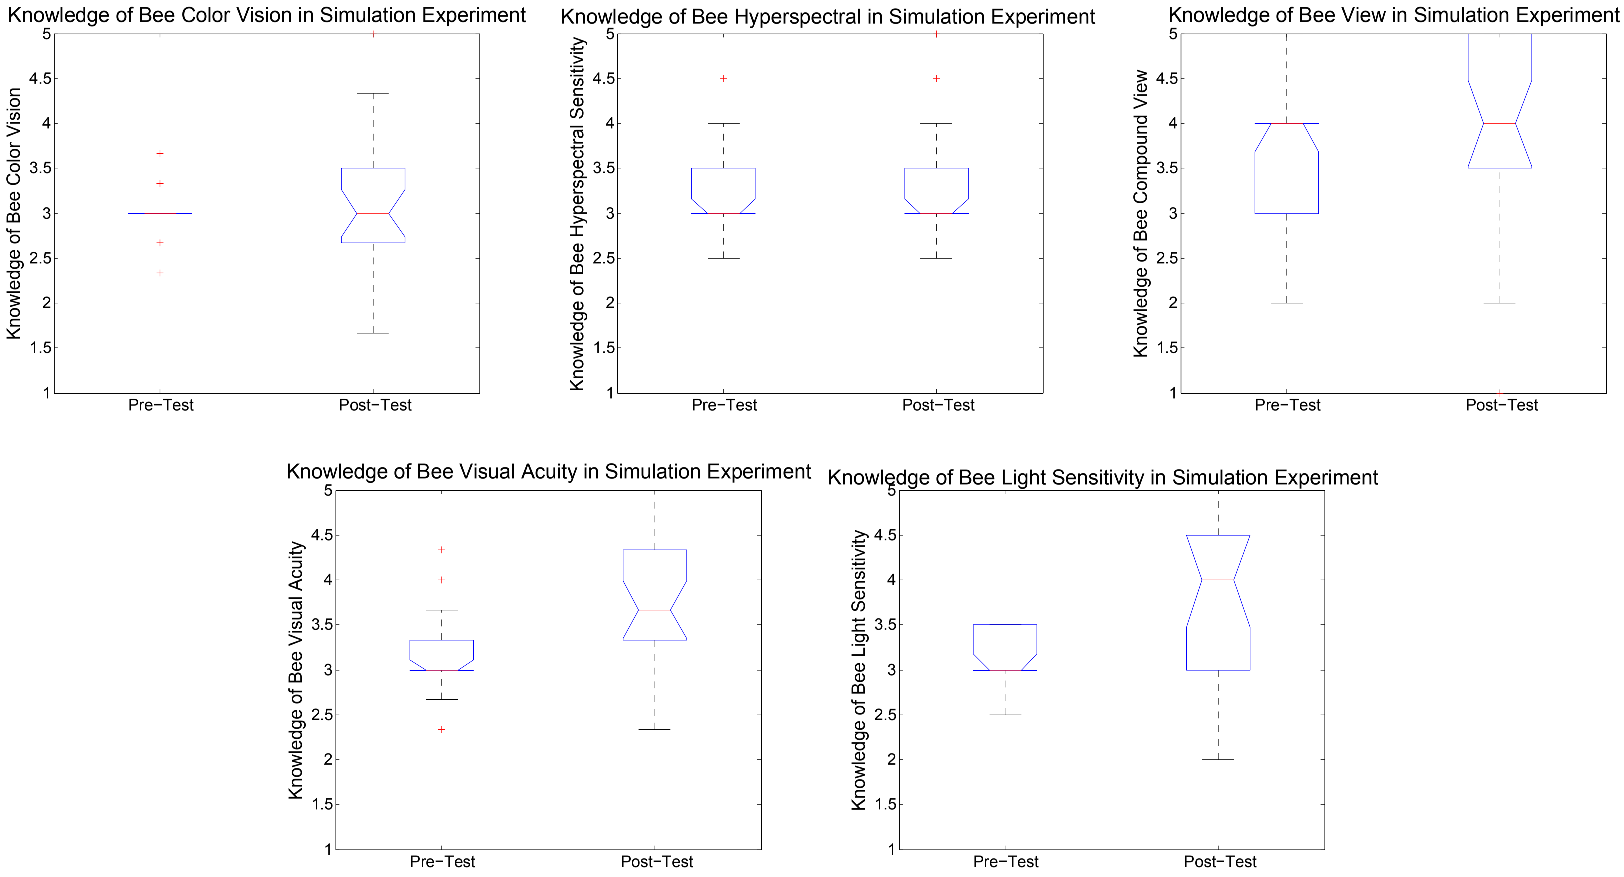

Figure 12.14: The mean results per visual characteristic on the pre- and post-tests for the 24 participants in the bee simulation experiments. The results are shown on a five point scale, where higher numbers indicate greater knowledge of the visual characteristic. Wilcoxon tests show statistically significant increases with regard to visual acuity and light sensitivity ($p < 0.05$).

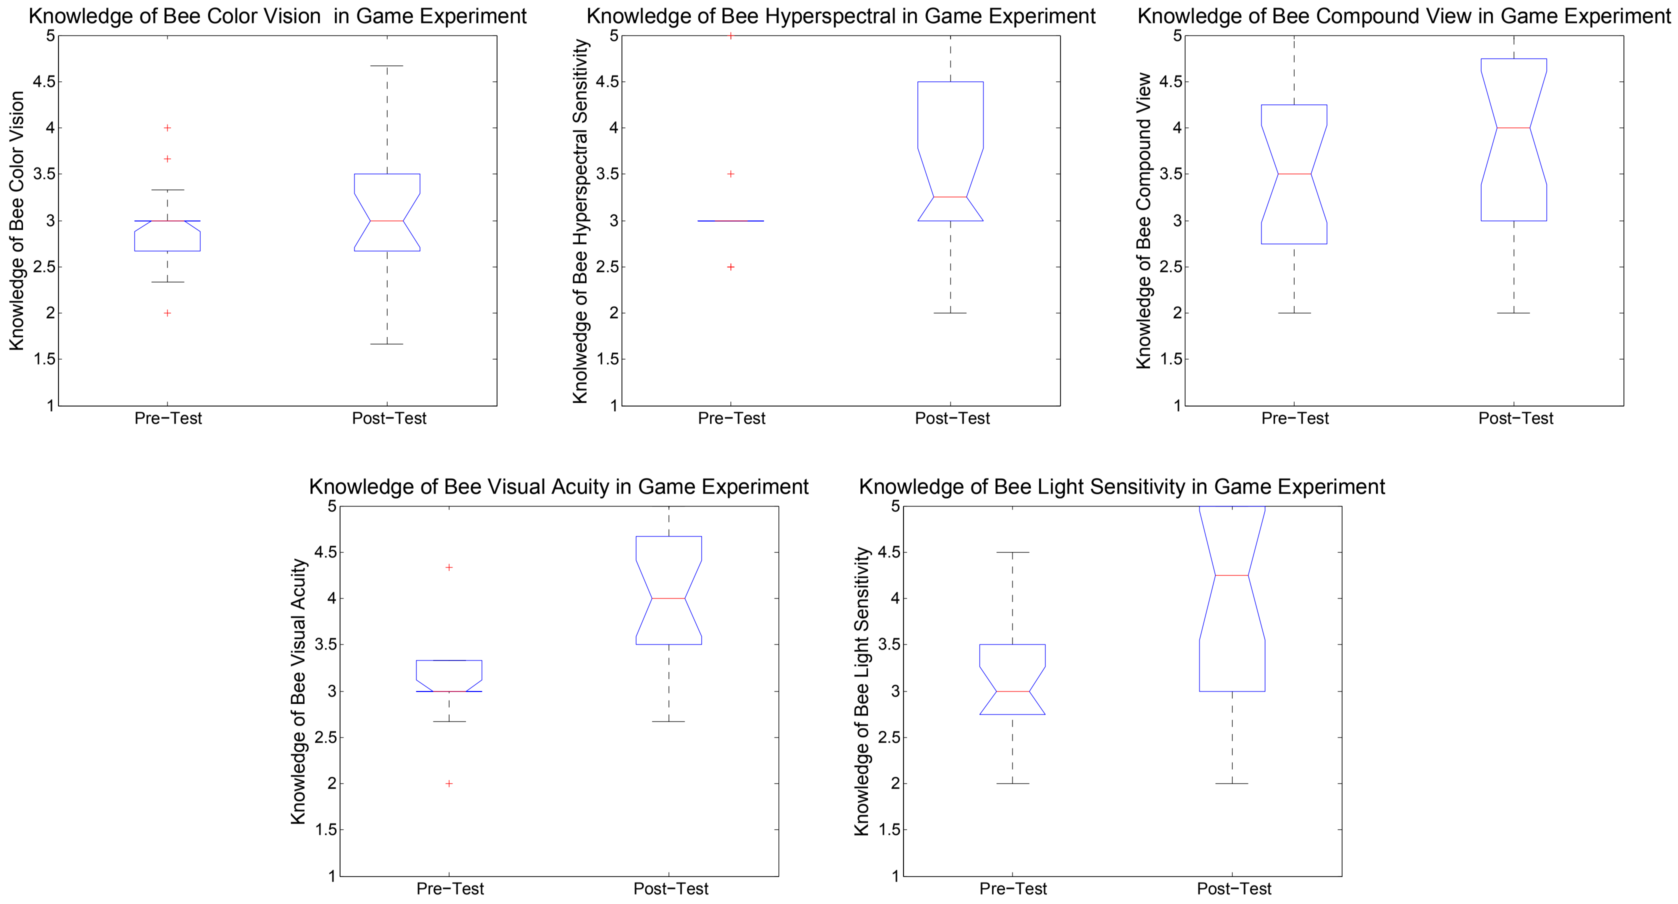

Figure 12.15: The mean results per visual characteristic on the pre- and post-tests for the 20 participants in the game experiment. The results are shown on a five point scale, where higher numbers indicate greater knowledge of the visual characteristics. Wilcoxon tests show statistically significant increases with regard to hyperspectral sensitivity, visual acuity, and light sensitivity ($p < 0.05$).

Another result worthy of discussion is the comparatively low learning increase realized in the category of color vision, as shown in Table 12.2. While designing the simulation, it was speculated that differences in color vision would be one of the easier facets to represent in a visual manner, and yet the results of this experiment indicate only marginal success in conveying this particular characteristic to observers. Interestingly, respondents were generally more capable of describing the relationship between bee and human color vision when responding to the short answer questions at the end of the experiment. This suggests that the wording used for the Likert items dealing with color vision may have been misinterpreted, which could have contributed to the result.

All the characteristics but the compound view seemed to be better conveyed when presented in the game rather than as a stand-alone simulation. It is possible that the game players were overloaded with other sensory data even as they engaged in a time-sensitive task, and as such did not have the opportunity to reflect on the nature of bees' compound view. Given that the difference between the learning increase seen between the two groups in this case is not statistically significant according to a Mann-Whitney test, it does not seem to be a basis for concern.

The participants in the game experiment were also asked to respond to Likert items dealing with their interest in bee and animal vision, and their perceived knowledge in this area both before and after playing *Bee Prepared*. Figure 12.16 displays the change in interest that occurred between the pre- and post-tests (M = 0.5, median = 0.5, s = 0.61). A Wilcoxon test suggests a statistically significant difference ($T = 4, p < 0.05$) between the results of these two tests, suggesting that the participants became significantly more interested and perceived themselves as more competent in the area of bee vision after playing the game. Games are generally expected to stimulate interest and engagement, so this result is not altogether surprising, but it is nonetheless worth noting that the players grew more interested in the underlying educational material present within the game.

At the end of the game experiment, respondents were asked to rate their enjoyment of the game on a scale of one to five. Figure 12.17 suggests that the responses were significantly positive (M = 4.15, median = 4, s = 0.5871), according to a binomial test with a test statistic (neutral point) of 3 ($p < 0.05$). While this is not a surprising result given that participants in such experiments tend to be predisposed towards them, it is nonetheless reassuring that the game was deemed enjoyable by first time participants, who at the same time demonstrated significant learning increases from the experience.

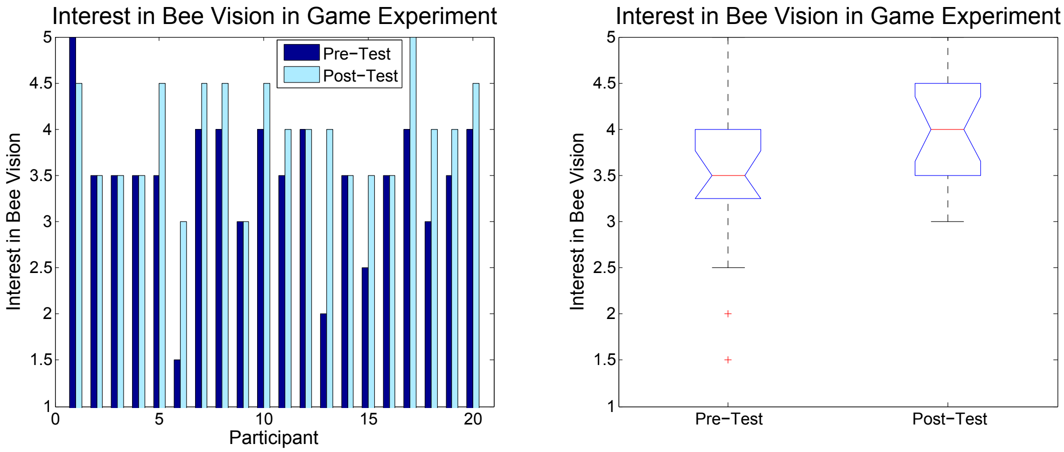

Figure 12.16: Mean interest in bee vision per participant recorded both before and after playing *Bee Prepared*. A Wilcoxon test suggests a statistically significant difference ($p < 0.05$) between the results of the pre- and post-tests, suggesting that the participants became significantly more interested in learning more about bee vision after playing the game.

Gameplay data was also recorded as part of the experiment. Without any extra assistance, 18 of the 20 respondents made it to the third mission or beyond during the fifteen minutes of play that they were allotted. This suggests that the instructions are generally sufficiently clear to allow the player to progress without any outside interference. Table 12.3 shows a correlation matrix that relates some of the game data that was logged for each participant with the percentage of post-test questions that the participants answered correctly for each of the five major visual characteristics that make up the bee simulation.

There are several statistically significant positive correlations shown in bold in Table 12.3. First of all, there is a significant positive correlation between computer and video game experience and game level achieved. Unsurprisingly, those accustomed to playing games and the more computer literate are more likely to succeed in the game. There is also a significant positive correlation between the game level reached and the percentage of chain bonuses achieved, suggesting that achieving more chain bonuses is an important factor in doing better at the game. However, there is no significant correlation between getting far in the game and achieving a better learning outcome. Even though less experienced players are less likely to reach high levels, they are still about as likely to emerge with beneficial learning outcomes. This suggests they achieved a concentrated sample without needing to progress deep into the game, which fulfills one of the design goals.

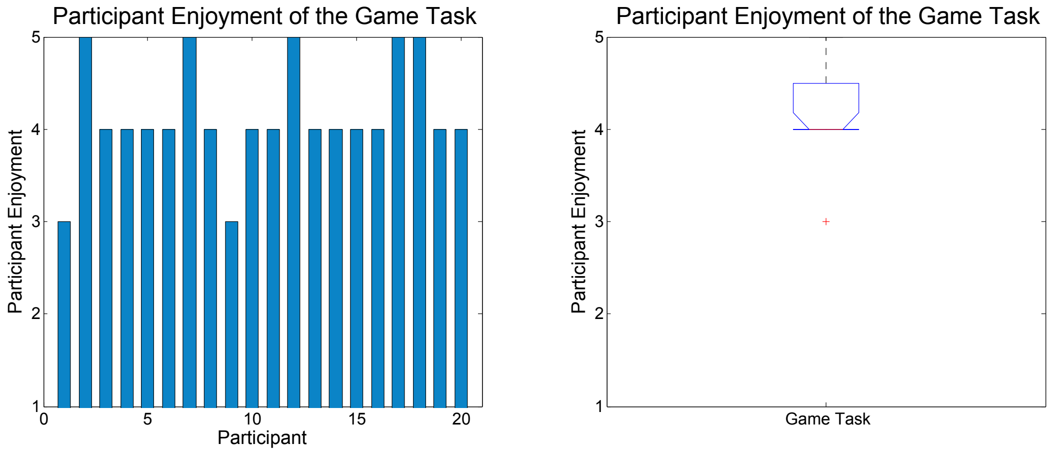

Figure 12.17: After playing *Bee Prepared*, players were asked to rate their enjoyment of the game on a scale from one to five. The results are statistically significant, according to a binomial test with test statistic of 3 ($p < 0.05$), and indicate that the participants found at least some part of the gameplay experience enjoyable.

There are two significant positive correlations recorded with regard to the learning outcomes. The first is between the percentage of chain bonuses and the visual acuity transformation. As players got a higher a percentage of chain bonuses, they also tended to have a better learning outcome in this area. This suggests that the process of obtaining more chain bonuses (searching the game world for flowers of similar colors) might have been beneficial in helping the player recognize the bee visual acuity transformation, and how blurry bee vision makes it difficult to see over longer distances. There is also a significant positive correlation between the level achieved and the test score of the respondents for the light sensitivity characteristic. Players who did better in the game tended to do better at identifying the bees' poor night vision. It's possible that these players took this characteristic into account while developing their gameplay strategies, and were thus able to address it more effectively.

The results of the *Bee Prepared* evaluation also suggest that more time should be spent constructing and evaluating the test instruments themselves to ensure that they will be interpreted in the manner in which they are intended. It may also be fruitful to consider the evaluation process during the design phase of an educational game, just as the educational content was considered when designing the game's core mechanics. Such an approach may facilitate harvesting in-game data that could be correlated with the results of other evaluation mechanisms to paint a more complete picture of a player's learning experience.

Correlation	Game Experience	Game Level	Chain Bonuses	Color Vision	Hyperspectral	Visual Acuity	Compound View	Light Sensitivity
Game Experience	1.0							
Game Level	**0.581**	1.0						
Chain Bonuses	0.246	**0.499**	1.0					
Color Vision	-0.024	0.046	-0.063	1.0				
Hyperspectral	0.018	0.092	0.073	0.052	1.0			
Visual Acuity	-0.140	0.211	**0.579**	0.253	0.214	1.0		
Compound View	0.188	0.172	-0.072	0.046	0.111	0.207	1.0	
Luminance	0.207	**0.558**	0.204	-0.102	-0.015	0.140	**0.497**	1.0

Table 12.3: The correlation coefficient matrix that considers how the participants' gameplay data relate to the percentage of correct responses they selected on the post-test for each of the five visual characteristics being represented by the bee simulation. Game experience includes the responses to several Likert items dealing with the participants' familiarity with computers and video games. Game level indicates the level in the game the participant reached during the fifteen minutes of play they were allotted. Chain bonuses refers to the percentage of seed pollinations where the seed color being gathered was the same as the last one the player pollinated, giving them a bonus. The cells in bold indicate significant correlations ($p < 0.05$).

Chapter 13

Conclusion

Throughout this dissertation, I have shown that the VDS methodology and framework is an effective way to produce illustrative simulations of visual systems that are capable of engaging and educating unprimed observers regarding animal vision. The framework uses physiological data about animals as input to generate simulations that can represent and emphasize the major differences between human vision and the visual systems being simulated.

The VDS methodology provides a means to represent visual characteristics that are outside the domain of human perception, such as hyperspectral sensitivity and motion tracking—characteristics that are an important part of perception for animals such as bees and pit vipers. However, there are still many visual characteristics that are challenging to represent in a meaningful manner, and it is an open question as to what sort of visual metaphor could successfully convey characteristics like tetrachromancy, metamerism, differences in depth perception, and oil drops in eyes, to name only a few. Nonetheless, the VDS framework demonstrates that visual representations can serve as a successful starting point to involve observers in learning more about animal vision and stimulating their interest in the topic.

It should be noted that the methodology employed to simulate animal visual systems is illustrative, and not intended to show an absolute representation of how an animal would see the world. This has been shown to be effective in stimulating interest and engagement in the topic, but it is important to remember that the illustrations do not represent absolute accuracy. Furthermore, the VDS framework does not necessarily represent the various visual characteristics with the visual salience they would possess within an animal's visual system. More examination could help determine the characteristics that scientists and educators believe should be emphasized as the most important elements of an animal's visual perception.

The game design principles that were employed to integrate the cat and bee simulations into educational games demonstrate the potential of serious games to enhance a visualization so that it conveys pertinent information more effectively, while at the same time engaging the player in an interactive and memorable experience. Although encouraging, these results underscore the importance of integrating educational material and evaluation mechanisms into all aspects of the game, especially the core game mechanics.

Evaluating the casual educational games documented in this dissertation remains challenging. The user studies described here drew participants from a unique population of adult students which may distort results and may not reflect the responses one might expect from the general population. Furthermore, the surveys used to gauge the participants' learning may not present an accurate view of the knowledge obtained from the games. Future educational games should strive to consider how the learning outcomes can be evaluated early in the design cycle, so that metrics can be incorporated into the game to gather data that can augment surveys as a means of assessing the player's learning experience.

Bibliography

[1] M. Agrawala and C. Stolte. Rendering effective route maps: improving usability through generalization. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 241–249, New York, NY, USA, 2001. ACM.

[2] L. W. Anderson, D. R. Krathwohl, P. W. Airasian, K. A. Cruikshank, R. E. Mayer, P. R. Pintrich, J. Raths, and M. C. Wittrock. *A Taxonomy for Learning, Teaching, and Assessing: A Revision of Bloom's Taxonomy of Educational Objectives, Abridged Edition.* Allyn & Bacon, 2 edition, 2000.

[3] K.C. Barnett, S.M. Crispin, A.G. Matthews, and J.D. Lavach. *Equine Ophthalmology.* Saunders, 2004.

[4] B. A. Barsky. Vision-realistic rendering: simulation of the scanned foveal image from wavefront data of human subjects. In *APGV'04*, pages 73–81, 2004.

[5] D. Bartz, D. W. Cunningham, J. Fischer, and C. Wallraven. State-of-the-art of the role of perception for computer graphics. In *Proceedings of the 29th Annual Conference Eurographics (EG 2008)*, pages 65–86, 2008.

[6] C.R. Braekevelt. Fine structure of the feline tapetum lucidum. *Anat Histrol Embryol*, 19:97–105, 1990.

[7] H. Brettel, F. Viénot, and J. D. Mollon. Computerized simulation of color appearance for dichromats. *J. Opt. Soc. Am. A*, 14(10):2647–2655, Oct 1997.

[8] K. Cater, A. Chalmers, and P. Ledda. Selective quality rendering by exploiting human inattentional blindness: looking but not seeing. In *Proceedings of the ACM symposium on Virtual reality software and technology*, VRST '02, pages 17–24, New York, NY, USA, 2002. ACM.

[9] M. Chi, T. Lee, Y. Qu, and T. Wong. Self-animating images: illusory motion using repeated asymmetric patterns. *ACM Trans. Graph.*, 27:62:1–62:8, August 2008.

[10] T. Christensen. *Methods in Insect Sensory Neuroscience.* CRC Press, 2004.

[11] Cognaxon. Snake eye vision. *http://www.cognaxon.com*, 2009.

[12] P. Coleman, K. Singh, L. Barrett, N. Sudarsanam, and C. Grimm. 3d screen-space widgets for non-linear projection. In *Proceedings of the 3rd international conference on Computer graphics and interactive techniques in Australasia and South East Asia*, GRAPHITE '05, pages 221–228, New York, NY, USA, 2005. ACM.

[13] M. F. Deering. A photon accurate model of the human eye. 24(3):649–658, July 2005.

[14] L. Doucet and V. Srinivasan. Designing entertaining educational games using procedural rhetoric: a case study. In *Sandbox '10: Proceedings of the 5th ACM SIGGRAPH Symposium on Video Games*, pages 5–10, New York, NY, USA, 2010. ACM.

[15] F. Durand. An invitation to discuss computer depiction. In *NPAR '02: Proceedings of the 2nd international symposium on Non-photorealistic animation and rendering*, pages 111–124, New York, NY, USA, 2002. ACM.

[16] A. G. Dyer. Bee discrimination of flower colours in natural settings. *Entomologia generalis*, 28:257–268, 2006.

[17] A. G. Dyer and S. K. Williams. Mechano-optical lens array to simulate insect vision photographically. *The Imaging Science Journal*, 53:209–213, 2005.

[18] R.V. Eck. Digital game-based learning: It's not just the digital natives who are restless. *EDUCAUSE Review*, 41(2):16–30, 2006.

[19] D. H. Foster, S.M.C. Nascimento, and K. Amano. Information limits on neural identification of coloured surfaces in natural scenes. *Visual Neuroscience*, 21:331–336, 2004.

[20] F.L. Fu, R.C. Su, and S.C. Yu. Egameflow: A scale to measure learners' enjoyment of e-learning games. *Comput. Educ.*, 52(1):101–112, 2009.

[21] T. Fullerton, L. M. Malamed, N. Sharkasi, and J. Vigil. Designing history: the path to participation nation. In *Sandbox '09: Proceedings of the 2009 ACM SIGGRAPH Symposium on Video Games*, pages 7–14, New York, NY, USA, 2009. ACM.

[22] J. P. Gee. What video games have to teach us about learning and literacy. *Comput. Entertain.*, 1(1):20–20, 2003.

[23] A. Giger. *B-Eye [electronic resource] : see the world through the eyes of a honeybee.* Andrew Giger, [Canberra] :, 1995.

[24] A. A. Gooch. *Preserving salience by maintaining perceptual differences for image creation and manipulation.* PhD thesis, Evanston, IL, USA, 2006. Adviser- J. Tumblin.

[25] A. A. Gooch, S. C. Olsen, J. Tumblin, and B. Gooch. Color2gray: salience-preserving color removal. In *SIGGRAPH '05: ACM SIGGRAPH 2005 Papers*, pages 634–639, New York, NY, USA, 2005. ACM.

[26] B. Gooch and A. Gooch. *Non-Photorealistic Rendering.* A. K. Peters, Ltd., Natick, MA, USA, 2001.

[27] B. Gooch, P. J. Sloan, A. Gooch, P. Shirley, and R. Riesenfeld. Interactive technical illustration. In *I3D '99: Proceedings of the 1999 symposium on Interactive 3D graphics*, pages 31–38, New York, NY, USA, 1999. ACM.

[28] V. I. Govardovskii, N. Fyhrquist, T. O. M. Reuter, D. G. Kuzmin, and K. Donner. In search of the visual pigment template. *Visual Neuroscience*, 17(4):509–528, 2000.

[29] E. Guenther and E. Zrenner. The spectral sensitivity of dark- and light-adapted cat retinal ganglion cells. *Journal of Neuroscience*, 13:1543–1550, 1993.

[30] D. Gutierrez. Kids wildly out of touch with natural world, bbc survey reveals. *Natural News.com (http://www.naturalnews.com/025064_natural_survey_BBC.html)*, 2008.

[31] M. P. J. Habgood. *The effective integration of digital games and learning content.* PhD thesis, 2007.

[32] M. Haller, C. Hanl, and J. Diephuis. Non-photorealistic rendering techniques for motion in computer games. *Comput. Entertain.*, 2:11–11, October 2004.

[33] S. Hartwell. How intelligent are cats? *http://www.messybeast.com/intelligence.htm*, 2004.

[34] B. Hoffman. The medium is not the only message. *Encyclopedia of Educational Technology. Retrieved May 24, 2011, from http://eet.sdsu.edu/eetwiki/index.php/The_medium_is_not_the_only_message*, 2009.

[35] M. Ichikawa, K. Tanaka, S. Kondo, K. Hiroshima, K. Ichikawa, S. Tanabe, and K. Fukami. Preliminary study on color modification for still images to realize barrier-free color vision. In *Systems, Man and Cybernetics, 2004 IEEE International Conference on*, volume 1, pages 36 – 41 vol.1, 2004.

[36] L. Jevbratt. Zoomorph - enabling interspecies collaboration. *Presented at ISEA, Belfast, UK, http://zoomorph.org*, 2009.

[37] R. Johns and J. Shaw. Real-time immersive design collaboration: conceptualising, prototyping and experiencing design ideas. *Journal of Design Research*, 5, 2006.

[38] Y. Kawagishi., K. Hatsuyama, and K. Kondo. Cartoon blur: nonphotorealistic motion blur. In *Computer Graphics International, 2003. Proceedings*, pages 276 – 281, 2003.

[39] K. Kiili and H. Ketamo. Exploring the learning mechanism in educational games. In *Information Technology Interfaces, 2007. ITI 2007. 29th International Conference on*, pages 357–362, June 2007.

[40] T.D. Lamb. Photoreceptor spectral sensitivities: Common shape in the long-wavelength region. *Vision Research*, 35(22):3083–3091, 1995.

[41] S. Lefebvre and F. Neyret. Pattern based procedural textures. In *Proceedings of the 2003 symposium on Interactive 3D graphics*, I3D '03, pages 203–212, New York, NY, USA, 2003. ACM.

[42] G. Lepouras and C. Vassilakis. Virtual museums for all: employing game technology for edutainment. *Virtual Real.*, 8(2):96–106, 2004.

[43] M. Livingstone. *Vision and Art: The Biology of Seeing*. Harry N. Abrams, 2002.

[44] J. Long, A. Estey, D. Bartle, S. Olsen, and A. A. Gooch. Catalyst: seeing through the eyes of a cat. In *FDG '10: Proceedings of the Fifth International Conference on the Foundations of Digital Games*, pages 116–123, New York, NY, USA, 2010. ACM.

[45] F. Lopez, R. Molla, and V. Sundstedt. Exploring peripheral lod change detections during interactive gaming tasks. In *Proceedings of the 7th Symposium on Applied Perception in Graphics and Visualization*, APGV '10, pages 73–80, New York, NY, USA, 2010. ACM.

[46] R. B. Lotto and D. Purves. The effects of color on brightness. *Natural Neurosci.*, 2(11):1010–1014, 1999.

[47] Y. Ma, X. Gu, and Y. Wang. Color discrimination enhancement for dichromats using self-organizing color transformation. *Inf. Sci.*, 179:830–843, March 2009.

[48] G. M. Machado, M. M. Oliveira, and L. A. F. Fernandes. A physiologically-based model for simulation of color vision deficiency. *IEEE Transactions on Visualization and Computer Graphics*, 15:1291–1298, 2009.

[49] R. L. Mandryk, M. S. Atkins, and K. M. Inkpen. A continuous and objective evaluation of emotional experience with interactive play environments. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 1027–1036, New York, NY, USA, 2006. ACM.

[50] M. Masuch, S. Schlechtweg, and R. Schulz. Speedlines: depicting motion in motionless pictures. In *ACM SIGGRAPH 99 Conference abstracts and applications*, SIGGRAPH '99, pages 277–, New York, NY, USA, 1999. ACM.

[51] R. Menzel, D.F. Ventura, H. Hertel, J. M. de Souza, and U. Greggers. Spectral sensitivity of photoreceptors in insect compound eyes: Comparison of species and methods. *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology*, 158:165–177, 1986. 10.1007/BF01338560.

[52] G. W. Meyer and D. P. Greenberg. Color-defective vision and computer graphics displays. *IEEE Comput. Graph. Appl.*, 8:28–40, September 1988.

[53] D. R. Michael and S. L. Chen. *Serious Games: Games That Educate, Train, and Inform.* Muska & Lipman/Premier-Trade, 2005.

[54] G. More and A. Burrow. Observing the learning curve of videogames in architectural design. In *IE '07: Proceedings of the 4th Australasian conference on Interactive entertainment*, pages 1–6, Melbourne, Australia, Australia, 2007. RMIT University.

[55] G. S. V. Mouat. Photoreceptor inputs to cat lateral geniculate nucleus cells. *Experimental Brain Research*, 59(2):242–248, 1985.

[56] K. Naden. Evaluation strategies for educational technology. *Federation of American Scientists*, 2007.

[57] H. O'Brien. Defining and measuring engagement in user experiences with technology. *PhD Dissertation. Dalhousie University, Halifax, NS.*, 2008.

[58] M. Pivec, O. Dziabenko, and I. Schinnerl. Aspects of game-based learning. In *Proceedings of I-KNOW 2003*, pages 216–225, 2003.

[59] M. Potmesil and I. Chakravarty. Modeling motion blur in computer-generated images. In *Proceedings of the 10th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '83, pages 389–399, New York, NY, USA, 1983. ACM.

[60] A. Raj and R. Rosenholtz. What your design looks like to peripheral vision. In *Proceedings of the 7th Symposium on Applied Perception in Graphics and Visualization*, APGV '10, pages 89–92, New York, NY, USA, 2010. ACM.

[61] J.M. Randel, B.A. Morris, C.D. Wetzel, and B.V. Whitehill. The effectiveness of games for educational purposes: A review of recent research. *Simulation & Gaming*, 23(3):261–276, 1992.

[62] K. Rasche, R. Geist, and J. Westall. Detail preserving reproduction of color images for monochromats and dichromats. *IEEE Comput. Graph. Appl.*, 25:22–30, May 2005.

[63] K. Rasche, R. Geist, and J. Westall. Re-coloring images for gamuts of lower dimension. *Computer Graphics Forum*, pages 423–432, 2005.

[64] K.E. Ricci, E. Salas, and J.A. Cannon-Bowers. Do computer-based games facilitate knowledge acquisition and retention? *Military Psychology*, 8(4):295–307, 1996.

[65] M. H. Rowe. Trichromatic color vision in primates. *News in physiological sciences an international journal of physiology produced jointly by the International Union of Physiological Sciences and the American Physiological Society*, 17(3):93–98, 2002.

[66] M. P. Salisbury, S. E. Anderson, R. Barzel, and D. H. Salesin. Interactive pen-and-ink illustration. In *SIGGRAPH '94: Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 101–108, New York, NY, USA, 1994. ACM.

[67] A. Secord. Weighted Voronoi stippling. In *Proceedings of the second international symposium on Non-photorealistic animation and rendering*, pages 37–43. ACM Press, 2002.

[68] R. P. Shuurmans and E. Zrenner. Responses of the blue sensitive cone system from the visual cortex and the arterially perfused eye in cat and monkey. *Vision Research*, 21:1611–1615, 1981.

[69] A. B. Sichert, P. Friedel, and J. L. van Hemmen. Snake's perspective on heat: Reconstruction of input using an imperfect detection system. *Phys. Rev. Lett.*, 97(6):068105, Aug 2006.

[70] A. J. Sillman, V. I. Govardovskii, P. Rohlich, J. A. Southard, and E. R. Loew. The photoreceptors and visual pigments of the garter snake (thamnophis sirtalis): a microspectrophotometric, scanning electron microscopic and immunocytochemical study. *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology*, 181:89–101, 1997. 10.1007/s003590050096.

[71] D. Simons and C. Chabris. Gorillas in our midst: Sustained inattentional blindness for dynamic events. *British Journal of Developmental Psychology*, 13:113–142, 1995.

[72] P. Stanicek. Color scheme designer. *http://colorschemedesigner.com*, 2002.

[73] D.G. Stavenga, R.P. Smits, and B.J. Hoenders. Simple exponential functions describing the absorbance bands of visual pigment spectra. *Vision Research*, 33(8):1011–1017, 1993.

[74] A. Stockman and L. T. Sharpe. The spectral sensitivities of the middle-and long-wavelength-sensitive cones derived from measurements in observers of known genotype. *Vision Research*, 40(13):1711–1737, 2000.

[75] A. Stockman, L. T. Sharpe, and C. Fach. The spectral sensitivity of the human short-wavelength sensitive cones derived from thresholds and color matches. *Vision Res*, 39(17):2901–27, 1999.

[76] N. Sudarsanam, C. Grimm, and K. Singh. Non-linear perspective widgets for creating multiple-view images. In *Proceedings of the 6th international symposium on Non-photorealistic animation and rendering*, NPAR '08, pages 69–77, New York, NY, USA, 2008. ACM.

[77] V. Sundstedt, E. Stavrakis, M. Wimmer, and E. Reinhard. A psychophysical study of fixation behavior in a computer game. In *Proceedings of the 5th symposium on Applied perception in graphics and visualization*, APGV '08, pages 43–50, New York, NY, USA, 2008. ACM.

[78] L. G. Tateosian, C. G. Healey, and J. T. Enns. Engaging viewers through non-photorealistic visualizations. In *NPAR '07: Proceedings of the 5th international symposium on Non-photorealistic animation and rendering*, pages 93–102, New York, NY, USA, 2007. ACM.

[79] M.P.S To, I.D. Gilchrist, T. Troscianko, J.S.B. Kho, and D. J. Tolhurst. Perception of differences in natural-image stimuli: Why is peripheral viewing poorer than foveal? *ACM Trans. Appl. Percept.*, 6:26:1–26:9, October 2009.

[80] J. Walraven and J. W. Alferdinck. Color displays for the color blind. In *Proc. IS&T and SID 5th Color Imaging Conf.*, pages 17–22. Soc. for Imaging Science and Technology, 1997.

[81] P. Willemsen and A. A. Gooch. Perceived egocentric distances in real, image-based, and traditional virtual environments. *Virtual Reality Conference, IEEE*, 0:275, 2002.

[82] S. Williams and A.G. Dyer. A photographic simulation of insect vision. *Journal of Ophthalmic Photography*, 20(1):10–14, 2007.

[83] S. K. Williams and A. G. Dyer. Mechano-optical lens array to simulate insect vision photographically. *The Journal of BioCommunication*, 34(1):E3–E7, 2008.

[84] G. Wyszecki and W. S. Stiles. *Color Science: Concepts and Methods, Quantitative Data and Formulae (Wiley Series in Pure and Applied Optics).* Wiley-Interscience, 2 edition, August 2000.

[85] F.W.M. Yip and A.C.M. Kwan. Online vocabulary games as a tool for teaching and learning english vocabulary. *Educational Media International*, 43(3):233–249, 2006.

[86] A. Yoshida, V. Blanz, K. Myszkowski, and H. Seidel. Perceptual evaluation of tone mapping operators with real-world scenes. In *Human Vision and Electronic Imaging X, SPIE*, 2005.

[87] T.D. Zuk. *Visualizing Uncertainty.* PhD thesis, Calgary, Alberta, Canada, 2008. Adviser-Sheelagh Carpendale.

Appendix A

Parameters used for Cat, Bee, and Pit Viper Simulations

This appendix documents the configuration files used as input to the VDS framework to generate the cat, bee, and pit viper simulations shown throughout this dissertation.

A.1 Cat Simulation

```
// Models cat color vision.
Color cone: 450
Color cone: 550
Saturation: 0.7

// Models cat motion sensitivity.
Motion detection: 0.7

// Models cat light sensitivity.
Light Sensitivity: Log 8.0

// Models cat visual acuity.
Acuity: 0.27

// Models cat field of view.
FoV: 220
```

A.2 Bee Simulation

```
Color cone: 360
Color cone: 450
Color cone: 520
Saturation: 0.7

// Models bee hyperspectral sensitivity.
UV

Light Sensitivity: 0.7

Acuity: 0.4

// Models bee compound eye construction.
// Specifies: the extent of the screen that will be tiled with
// hexagonal viewports, the dimensions of the viewports,
// spacing between them, and their overlap.
HEXAGONS -1.1, -1.1, 1.1, 1.1, 0.05, 0.003, 0.66
```

A.3 Pit Viper Simulation

```
Color Cone: 430
Color Cone: 550
Saturation: 0.7

Acuity: 0.4

// Models pit viper motion sensitivity.
Motion detection: 0.95

// Models pit viper hyperspectral sensitivity.
IR

Light Sensitivity: Log 6.0
```

Appendix B

Calculating the Color Transformation Matrix

Given discretized response curves for humans and some other visual system, the goal is to devise a new set of cone response curves that contain only radiance distributions that are visible to both humans and the visual system being simulated.

B.1 Definitions

Let H be a matrix that stores discretized human color response curves. Thus, if $\omega(\lambda)$ is some radiance distribution, discretized as $\mathbf{x}$, and $R(\lambda)$, $G(\lambda)$, $B(\lambda)$ are the continuous red, blue, and green response functions:

$$Hx \approx \begin{pmatrix} \int R(\lambda)\omega(\lambda)d\lambda \\ \int G(\lambda)\omega(\lambda)d\lambda \\ \int B(\lambda)\omega(\lambda)d\lambda \end{pmatrix}.$$

Define C similarly, using some other visual system's response curves.

B.2 "Invisible" Light

If $C\mathbf{z} = 0$, then the radiance distribution $\mathbf{z}$ is *invisible* to the visual system that we are seeking to simulate. Any radiance distribution $\mathbf{x}$ can be altered by adding or subtracting multiples of such invisible $\mathbf{z}$ vectors, and the color perceived by that other visual system will remain unchanged. Vectors $\mathbf{z}$ having the property $C\mathbf{z} = 0$ are said to belong to the *null space* of C.

The aim is to create H', a modified version of the human response functions stored in H, with the property that radiance distributions invisible to the visual system being simulated are also invisible to the response curves of the modified human. However, insomuch as it is possible, the colors perceived by the modified human should be identical to the colors seen by the unmodified human. Formally, if N denotes the null space of C, our conditions are,

$$H'\mathbf{x} = 0 \quad \text{if} \quad \mathbf{x} \in N.$$

$$H'\mathbf{x} = H\mathbf{x} \quad \text{if} \quad \mathbf{x} \text{ is orthogonal to N.}$$

These two conditions are sufficient to uniquely specify H'. The space containing all $\mathbf{x}$ orthogonal to N is the *row space* of C. By creating a matrix M which projects input vectors into the row space of C, then maps them back to the domain of C, we can form the unique matrix M with the properties that,

$$M\mathbf{x} = 0 \quad \text{if} \quad \mathbf{x} \in N.$$

$$M\mathbf{x} = \mathbf{x} \quad \text{if} \quad \mathbf{x} \text{ is orthogonal to N.}$$

Defining $H' := HM$ creates an H' with the desired properties.

B.3 Calculating M

There are many equivalent ways of deriving the matrix that will remove all components of a radiance distribution that are invisible to the visual system being simulated. Depending on the nature of one's background in linear algebra, some of these definitions may prove more intuitive than others. Here is a list of equivalent Matlab expressions that return M, along with how each one can be motivated:

`orth(C') * orth(C')'`: Projecting $\mathbf{x}$ into the row-space of C, then mapping the result back to the domain.

`eye(size(C,2)) - null(C) * null(C)'`: Subtracting out the null space component of $\mathbf{x}$.

`pinv(C) * C`: Given a distribution $\mathbf{x}$, first find the colors for the system being simulated $\mathbf{r} = C\mathbf{x}$, then apply the pseudo-inverse of C to $\mathbf{r}$, thus finding a radiance distribution $\hat{\mathbf{x}}$ having the property that $C\hat{\mathbf{x}} = \mathbf{r}$.

Because the pseudo-inverse always returns an $\hat{\mathbf{x}}$ inside the row space of C, and because the row space and null space of a matrix are guaranteed to form a mutually orthogonal partitioning of the domain of C, all of the above matrices are identical.

B.4 Approximating Radiance Distributions

If the light being used as input to the color transformation was defined in terms of radiance distributions, changing H to H' would be sufficient. However, for the sake of convenience, it is desirable to be able to accept input data that specifies light using the conventional combination of red, green, and blue values, as this is the form of input that will come from a standard virtual environment or game engine. Therefore, given three color values stored in the vector $\mathbf{c}$, a means of inferring a radiance distribution $\mathbf{x}(\mathbf{c})$ is required. Given such a function, colors seen by the modified human response curves can be calculated, $\mathbf{c}' := H' * \mathbf{x}(\mathbf{c})$.

One natural condition to put on $\mathbf{x}(\mathbf{c})$ is $H\mathbf{x}(\mathbf{c}) = \mathbf{c}$. However, there are a large number of different $\mathbf{x}(\mathbf{c})$ definitions that satisfy this constraint. This dissertation chooses to infer $\mathbf{x}$ values using the psuedo-inverse of H. $\mathbf{x}(\mathbf{c})$ is defined to be $P\mathbf{c}$, where P is the pseudo-inverse of H. (Equivalently, one can say that $\mathbf{x}(\mathbf{c})$ is defined to be the value $\mathbf{x}$ for which $H\mathbf{x} = \mathbf{c}$ and $||\mathbf{x}||$ is as small as possible.)

This choice of $\mathbf{x}(\mathbf{c})$ has the advantage of allowing the mapping from $\mathbf{c}$ to $\mathbf{c}'$ to be expressed as a $3x3$ matrix. Specifically,

$$\mathbf{c}' = H'x(\mathbf{c}) = H'P\mathbf{c}.$$

Thus, $\mathbf{c}' = S\mathbf{c}$, where $S := H'P$. Given definitions for H and C, S can be calculated in the following one-line Matlab command:

```
S = H * pinv(C) * C * pinv(H);
```

In summary, given a vector (c) that stores the RGB color perceived by a human, this method can generate the color vector c' that will be perceived by the modified human. The vector c' is given by the transform,

$$\mathbf{c}' = S\mathbf{c}.$$

Achieving the color transformation as a single matrix multiply makes it easy to apply as a real time transformation on a video feed or a virtual environment as part of a pixel shader.

B.5 Metamerism

Forcing $\mathbf{x}$ to be a function of $\mathbf{c}$ has one potentially significant drawback. It implies that one cannot represent materials that are metamers to human vision, but distinguishable colors for the visual system being simulated [19]. It is not clear how common such materials are in the real world, nor how they could be intuitively represented as part of a color transformation.

B.6 Color Discrimination

When creating content intended to highlight the differences between human color discrimination and that of the visual system being simulated, it is useful to know how one can change the color perceived by a human without changing the color perceived by the system being simulated.

Let $\mathbf{v}$ be a vector that can be added to an existing radiance distribution in order to change the color that will be perceived by a human, without changing the color perceived by the visual system being simulated. The space V containing all $\mathbf{v}$ with the property that $H\mathbf{v} \neq 0$ and $C\mathbf{v} = 0$ is very large. There is, however, a one dimensional subspace, $\hat{V} \subset V$ with particular importance. Namely, $\hat{V}$ defined so that all $\hat{\mathbf{v}} \in \hat{V}$ are both an element of the null space of C and the row space of H.

The fact that $\hat{\mathbf{v}}$ lies inside the row-space of H is important, as the mapping of color to radiance distributions, $\mathbf{x}(\mathbf{c}) := P\mathbf{c}$, will always return $\mathbf{x}$ values inside the row-space of H. Therefore, the only way to alter a color $\mathbf{c}$ without changing the implied $\mathbf{c}' = S\mathbf{c}$ is to change $\mathbf{c}$ to be $\mathbf{c} = \mathbf{c} + aH\hat{\mathbf{v}}$, where a is any scalar constant.

In Matlab, a basis for $\hat{V}$ can be derived, again, using a simple one line expression:

$$\hat{\mathtt{v}} = \mathtt{H}' * \mathtt{null}(\mathtt{C} * \mathtt{H}');$$

Appendix C

Evaluation of Cat Simulation

Upon giving consent to participate in the Catalyst evaluation, participants were immediately given the instructional text shown in Figure C.1.

C.1 Pre-Test

After reading and responding to the instructional document, participants were next presented with a pre-test that included seven five-point Likert items. The Likert items used in the experiments measured the subject's level of agreement, with the responses ranging from strongly disagree to strongly agree. The items were displayed by a computer program, as shown in Figure C.2, in a randomized order for each participant.

- Likert items measuring video game experience
  - I play a lot of video games.
  - I am not very good at video games.
  - I am interested in video games.
- Likert items measuring interest in cats and their vision
  - I am not interested in learning more about the cat visual system.
  - I would feel confident sharing my knowledge of cat vision with someone who is interested.
  - I know more about the cat visual system than the average person.
  - I do not find animals interesting.

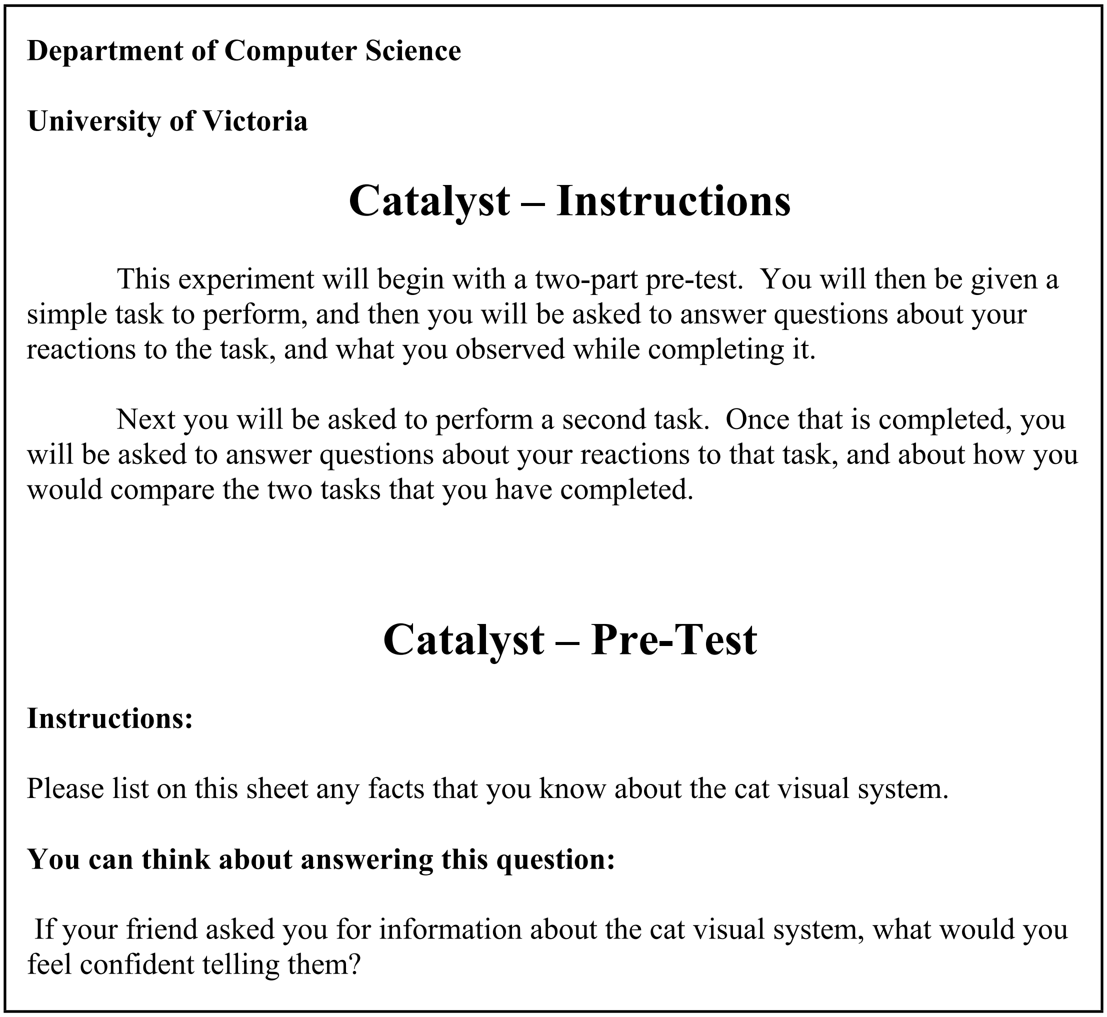

Department of Computer Science

University of Victoria

Catalyst – Instructions

This experiment will begin with a two-part pre-test. You will then be given a simple task to perform, and then you will be asked to answer questions about your reactions to the task, and what you observed while completing it.

Next you will be asked to perform a second task. Once that is completed, you will be asked to answer questions about your reactions to that task, and about how you would compare the two tasks that you have completed.

Catalyst – Pre-Test

Instructions:

Please list on this sheet any facts that you know about the cat visual system.

You can think about answering this question:

If your friend asked you for information about the cat visual system, what would you feel confident telling them?

Figure C.1: The instructional text that was given to particiants in the *Catalyst* evaluation. It outlines the general experimental procedure and gives respondents an unstructured opportunity to enumerate any knowledge they have regarding cat vision. These responses are used to determine their a priori knowledge going into the experiment.

C.2 Post-Test

After completing the learning task, pariticpants were given the freeform questionaire shown in Figure C.3, where they could enumerate any additional knowledge they had acquired about cat vision through participating in the learning activity. Upon completing this form, they were presented with the second part of the post-test, which included the Likert items from the pre-test dealing with interest in cats and their vision, in addition to Likert items measuring their level of engagement while completing the learning task.

- Likert items measuring engagement
    - I was really drawn into this learning task.
    - I felt involved in this learning task.
    - This learning experience was fun.
    - This learning task was worthwhile.
    - I consider this learning experience a success.
    - This learning experience did not work out the way I had planned.
    - This learning experience was rewarding.
    - I would recommend this learning experience to my friends and family.

They were then tasked with reading the educational document shown in Figure C.4, which contains information about cats and their vision.

C.3 Second Post-Test

After reading the text document, the participants were subjected to another post-test that included the engagement items shown in Section C.2, in addition to several new items that asked them to directly compare the first learning experience against reading the educational document shown in Figure C.4.

- Likert items comparing *Catalyst* and the text reading task
    - I feel that the text reading task was more educational than the other task.
    - I enjoyed the text reading task more than the other task.
    - I was less interested in the text reading task than the other task.
    - The text reading task did not make me think as much as the other task.

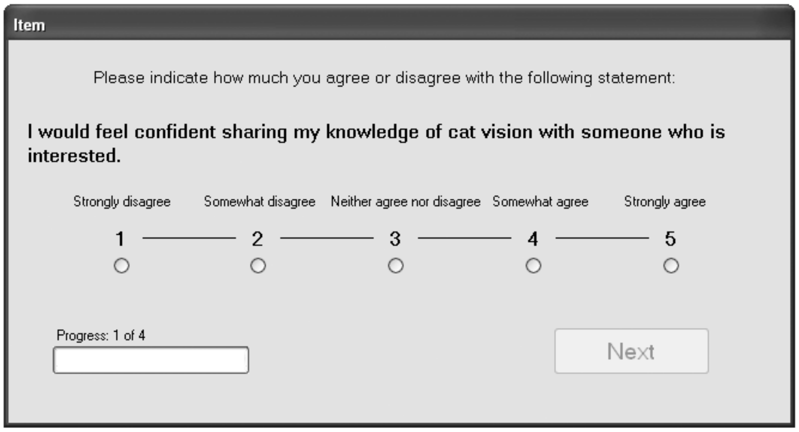

Item

Please indicate how much you agree or disagree with the following statement:

I would feel confident sharing my knowledge of cat vision with someone who is interested.

Strongly disagree	Somewhat disagree	Neither agree nor disagree	Somewhat agree	Strongly agree
1	2	3	4	5
○	○	○	○	○

Progress: 1 of 4

Figure C.2: A Likert item used in the *Catalyst* evaluation. All Likert items described in this document used the standard five-point responses shown here. The Likert items were presented to the participants by a computer program in a randomized order, and half the questions were reverse-coded.

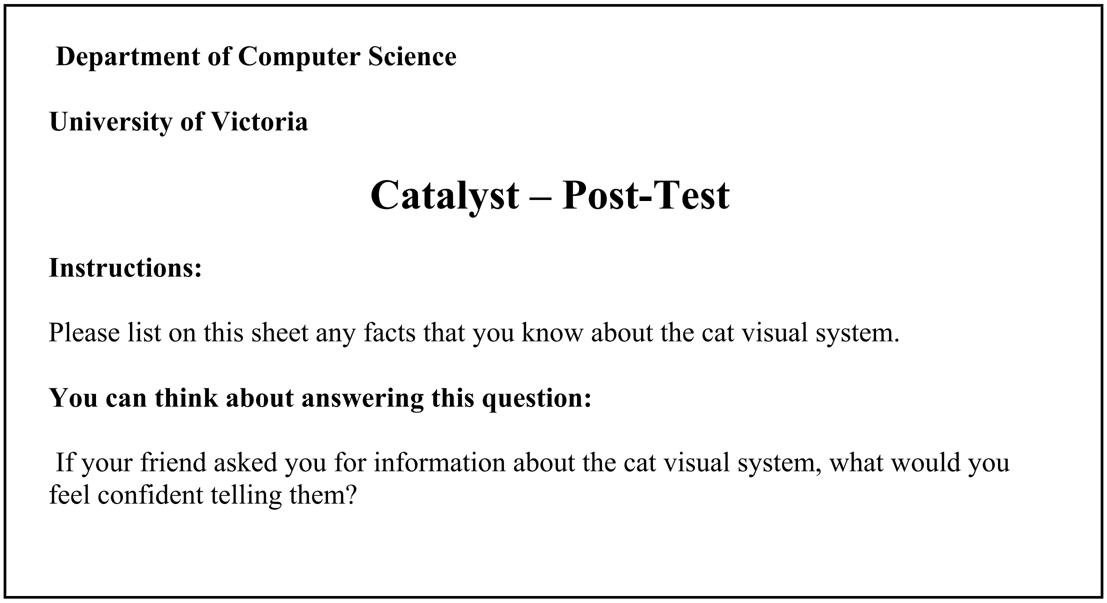

Department of Computer Science

University of Victoria

Catalyst – Post-Test

Instructions:

Please list on this sheet any facts that you know about the cat visual system.

You can think about answering this question:

If your friend asked you for information about the cat visual system, what would you feel confident telling them?

Figure C.3: After the first learning experience, participants were given another form to recount their observations regarding cat vision.

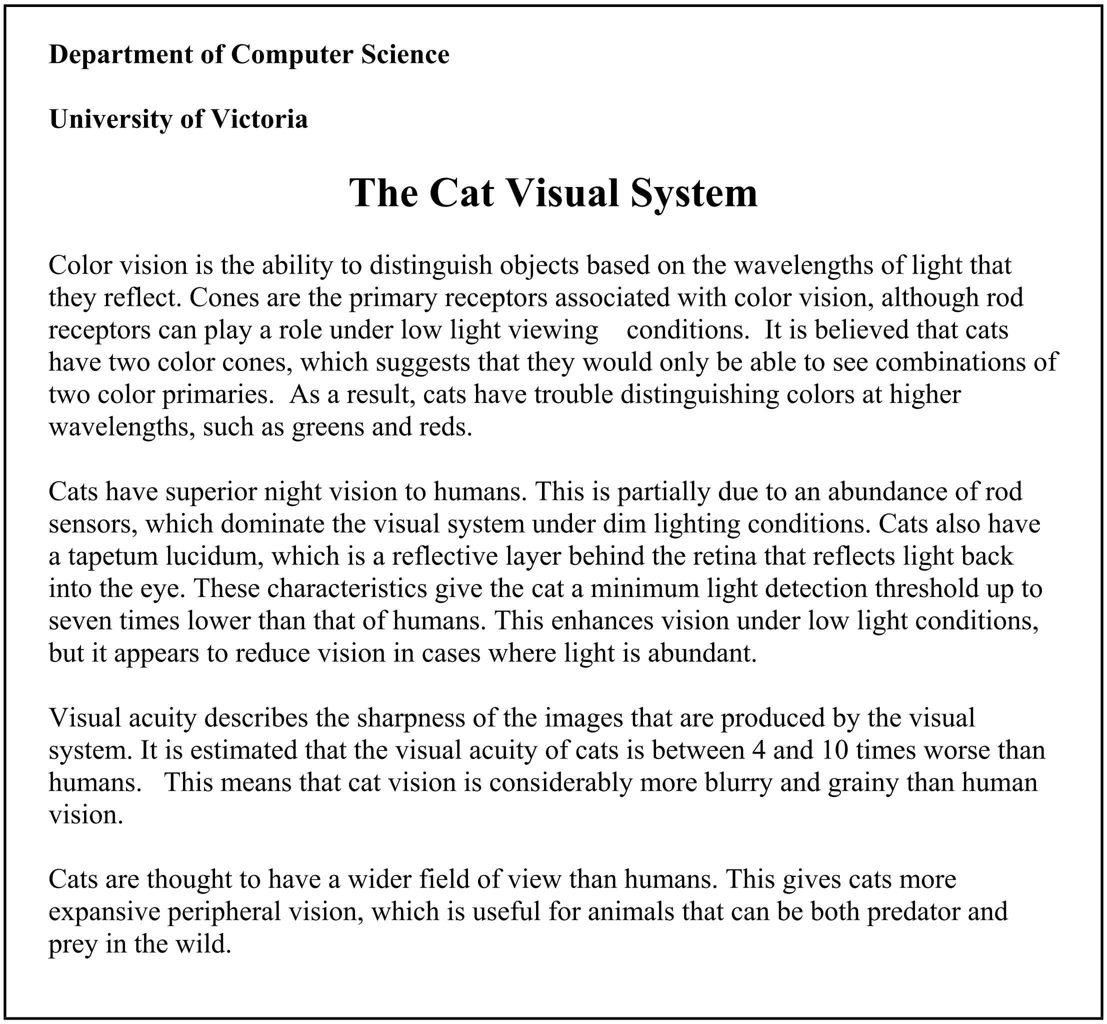

Department of Computer Science

University of Victoria

The Cat Visual System

Color vision is the ability to distinguish objects based on the wavelengths of light that they reflect. Cones are the primary receptors associated with color vision, although rod receptors can play a role under low light viewing conditions. It is believed that cats have two color cones, which suggests that they would only be able to see combinations of two color primaries. As a result, cats have trouble distinguishing colors at higher wavelengths, such as greens and reds.

Cats have superior night vision to humans. This is partially due to an abundance of rod sensors, which dominate the visual system under dim lighting conditions. Cats also have a tapetum lucidum, which is a reflective layer behind the retina that reflects light back into the eye. These characteristics give the cat a minimum light detection threshold up to seven times lower than that of humans. This enhances vision under low light conditions, but it appears to reduce vision in cases where light is abundant.

Visual acuity describes the sharpness of the images that are produced by the visual system. It is estimated that the visual acuity of cats is between 4 and 10 times worse than humans. This means that cat vision is considerably more blurry and grainy than human vision.

Cats are thought to have a wider field of view than humans. This gives cats more expansive peripheral vision, which is useful for animals that can be both predator and prey in the wild.

Figure C.4: Text containing educational material related to cats and their visual system. Participants in both conditions were asked to read this after their first learning experience, and to compare the two as part of the second post-test.

Appendix D

Evaluation of Bee Simulation

D.1 Standard Test Instrument

The standard testing instrument shown here was used in all three phases of the bee evaluation, and included the five-point Likert items listed below. All the Likert items measured the participants' agreement with correct and incorrect statements about bee vision. Participants' knowledge about a characteristic of bee vision was determined by how strongly they agreed with correct statements and disagreed with incorrect ones. The items were displayed by a computer program, as shown in Figure C.2, in a randomized order for each participant.

- Color vision: the responses to the following Likert items were averaged to determine the subject's knowledge of bee color vision.
  - Humans are better at seeing blue colors than bees.
  - Humans are better at seeing green colors than bees.
  - Humans are better at seeing red colors than bees.

- Hyperspectral Sensitivity: the responses to the following Likert items were averaged to determine the subject's knowledge of bee hyperspectral sensitivity.
  - Bees can perceive ultraviolet light.
  - Bees can perceive infrared light.

- Compound Eye Construction: the responses to the following Likert items were averaged to determine the subject's knowledge of bee compound eye construction.

  - Human see a single view of a scene at once.
  - Bee see many different views of a scene at once.

- Visual acuity: the responses to the following Likert items were averaged to determine the subject's knowledge of bee visual acuity.
  - Bees are better at seeing over long distances than humans.
  - Bees are better at seeing things that are near than humans.
  - Bees are better able to focus their vision than humans.
- Light Sensitivity: the responses to the following Likert items were averaged to determine the subject's knowledge of bee light sensitivity.
  - Bees can see better in the dark than humans.
  - Bee vision gets more blurry at night.

D.2 Performance Phase

Several Likert items were also included as part of the post-test for the performance phase of the experiment. After completing the color counting task under both the glow and false coloring conditions, participants responded to the following Likert items, which directly compared the two conditions in terms of aesthetics, perceived learning outcomes, and personal preference.

- Likert items comparing glow and false coloring for representing hyperspectral sensitivity
  - The first task was more aesthetically pleasing than the second.
  - I feel that I learned more about bee vision during the first task than the second.
  - I preferred the second task to the first.

D.3 Game Phase

The game phase of the evaluation included several more Likert items specific to that part of the experiment. The pre-test included the following Likert items related to the player's computer and video game experience, in addition to items regarding their interest and perceived knowledge in bees and their vision.

- Likert items measuring computer and video game experience
    - I am very familiar with using computers.
    - I play a lot of video games.
    - I am not very good at video games.
- Likert items measuring interest in animal and bee vision
    - I am not interested in learning more about how animals see the world.
    - I know more about bee vision than the average person.

The post-test for the game phase of the evaluation also included the interest questions shown above, in addition to one asking the participant to assess their enjoyment while playing the game.

- Likert item measuring enjoyment of *Bee Prepared*
    - I enjoyed playing the game.