An Empirical Evaluation

of

Interactive Colour Selection Techniques

by

Michael Wilhelm Schwarz

A thesis
presented to the University of Waterloo
in partial fulfilment of the
requirements for the degree of
Master of Mathematics
in
Computer Science

Waterloo, Ontario, 1985

I hereby declare that I am the sole author of this thesis.

I authorize the University of Waterloo to lend this thesis to other institutions or individuals for the purpose of scholarly research.

Signature

I further authorize the University of Waterloo to reproduce this thesis by photocopying or other means, in total or in part, at the request of other institutions or individuals for the purpose of scholarly research.

Signature

The University of Waterloo requires the signatures of all persons using or photocopying this thesis. Please sign below, and give address and date.

An Empirical Evaluation of Interactive Colour Selection Techniques

Michael W. Schwarz

Computer Graphics Laboratory
Department of Computer Science
University of Waterloo
Waterloo, Ontario

ABSTRACT

The choice of colours used in an image has a profound effect on how the intended viewer perceives and interprets its contents. To have proper control over the colours used in creating an image, the user of a computer graphics system must be able to use an appropriate method for colour selection. Although a number of colour selection techniques have been developed, there is virtually no empirical evidence to suggest which of these techniques are well suited for interactive computer graphics applications or what factors are important for designing an effective user interface for colour selection. In order to investigate these issues, a comprehensive examination of the colour selection process has been undertaken. The initial phase of this study, a psychophysical experiment based on the method of colour matching, has been completed. In this experiment, a number of colour selection techniques are evaluated to determine their suitability for computer graphics applications. This thesis describes in detail the objectives, design considerations, implementation, results and implications of the colour matching experiment. In addition, a number of methods for performing colour selection tasks are discussed.

Acknowledgements

This thesis was prepared under the supervision of Dr. John C. Beatty using the typesetting facilities of the Computer Graphics Laboratory at the University of Waterloo. The readers were Dr. William B. Cowan and Dr. Bruce W. Char. The experiment described herein was developed under the direction of Dr. Cowan.

Credit is due to all those involved in this project, especially Dr. Beatty and Dr. Cowan who were primarily responsible for its conception and realization. Special recognition is also due to the contributions of Dr. Gunther Wyszecki who provided invaluable advice and encouragement and donated a substantial amount of his time to fulfil the role of "experienced subject" (under rather chaotic circumstances). The experiment was conducted with the reliable assistance of Dr. Colin Ware and Michael Morgan. Advice on statistical issues was provided by Dr. Jane F. Gentleman in the early stages of the experiment's design.

I'ld like to express my gratitude to all the people mentioned above, to all those who volunteered to participate in the experiment, to the members of the Computer Graphics Lab for their cooperation and patience, to my parents for much needed moral and financial support, to David Forsey and Ian! Allen for keeping me well supplied with Shreddies, Wheat Thins and 1200 baud modems, to Paul Philp for sending me many wonderful postcards from Europe, and especially to all my friends who made my stay in Waterloo a truly *educational* experience.

Contents

1. Introduction

To be an effective tool, an interactive computer graphics system must provide an environment in which the user's ability to accomplish or create can be enhanced [Baecker80]. The interaction that occurs between the system and its users should provide convenient access to the power and functionality of the system. Despite recent advances in the development of systematic methodologies for the design of human-computer interfaces, the task of constructing an effective user interface for many computer graphics applications remains a difficult process. Much work is still needed in many areas before the design of human-computer interfaces becomes a well defined science [Buxton83].

One area in which the design of the user interface is still particularly unsystematic and difficult is in applications where operations involving *colour* must be integrated into an interactive environment. The complex nature of the human visual system and the fact that the perception of colour is not completely understood severely complicate the design process. As computer systems with high quality colour capability become increasingly accessible, the need for practical design rules that can aid the system designer in incorporating sophisticated colour operations into the user interface is becoming more urgent.

This thesis deals at length with one user interface design issue that has implications for many applications of computer graphics - *colour selection*. The method by which the user specifies which colour (or colours) are to be used in subsequent operations has a considerable impact on a typical system's overall usability. This is particularly true in applications where the selection of colours must be frequently performed. Although a number of techniques for colour selection have been developed, there is virtually no empirical evidence to suggest which of these systems is best suited for interactive computer graphics applications. Neither is there any concrete information that indicates what factors are important for designing an effective user interface for colour selection tasks.

In order to investigate the task of colour selection in an extensive manner, a comprehensive examination of colour selection techniques and the issues that affect the performance of users executing colour manipulation operations has been undertaken. The initial phase of this study, a psychophysical experiment based on the method of colour matching, has been completed. The primary goal of this experiment was to evaluate empirically a number of colour selection techniques to determine their suitability for interactive computer graphics applications. The colour matching experiment is the main focus of this thesis. Its objectives, design considerations, implementation, results and implications are described here in detail.

The problem of colour selection is explicitly defined in Chapter 2. A discussion of some of the factors that make the process of colour selection difficult is given. In addition, colour editing, a task that is closely related to colour selection, is defined and briefly discussed. An overview of the relevant computer graphics literature is also provided in this chapter.

Chapter 3 describes the methodology used in this investigation to analyze the problem of colour selection. A discussion of the objectives that the colour matching experiment was designed to achieve is presented. The physiological and psychophysical methods of investigation are distinguished and compared and a description of the colour matching operation is given. The reasons for applying the psychophysical approach and, in particular, the task of colour matching to the problem of selecting colours are also discussed.

In Chapter 4, an overview of several methods for selecting colour in a computer graphics system is presented. Although emphasis is placed on selection techniques that are capable of specifying colours at arbitrary levels of precision, a number of methods that are more suitable for applications that require only a limited number of colours are also discussed. A detailed description of the twelve colour selection systems evaluated in the experiment, including a discussion of why these particular systems were chosen for analysis, is given here.

Chapter 5 gives a detailed description of the colour matching experiment. Topics of discussion include: the equipment used, the environment in which the experiment was conducted, the exact procedure subjects were required to follow, how subjects were selected and trained and what data was collected while subjects performed the experiment.

The results of the experiment are presented in Chapter 6. A number of criteria are used to measure the relative effectiveness of the colour selection systems, including: the closeness of the final match, how long it took to achieve the match and the time taken to cross various "thresholds of colour difference."

From these results, conclusions are drawn about the suitability of the evaluated selection systems for use in interactive computer graphics environments. These conclusions are presented in Chapter 7. In addition, recommendations are made about the design of future colour selection systems.

Chapter 8 proposes a number of possible extensions to the work begun here. Suggestions for further analysis of the data generated from this experiment, as well as proposals for future research projects that are related to this one are made.

2. Problem Description

As the use of computer graphics technology has become more prevalent, the level of sophistication of computerized graphical applications has increased dramatically. Interactive picture creation systems, for example, with which users can interactively create and edit complex digital images by manipulating an interaction device such as a tablet or a mouse, typify this recent surge in functionality [Smith82, Plebon82]. Such systems are now frequently utilized in a variety of complex applications including document preparation, commercial animation production, and computer generated art.

The *users* of computer graphics systems have also become substantially more sophisticated. For instance many users of computer-driven picture creation systems, such as graphic artists and animation specialists, are professionals in their own fields and as a result have complicated user needs and relatively strong performance expectations. Hence, not only must a system perform an abundance of functions and be convenient to use, but it must also be capable of producing high quality images that are "visually appealing." Regardless of whether this means that the user is more concerned about the informative aspects of the images created (i.e. how information contained in an image is interpreted) or their artistic merit, the system must be able to efficiently produce images at a high enough level of aesthetic quality that meet the application's needs [Trucken81, Goetz82].

Any application that requires the production of visually appealing graphical images needs to use *colour* effectively. The choice of colours used, both separately and in combination, has a profound effect on how the intended viewer perceives and interprets the contents of the image. Used effectively, colour can complement and even enhance the information content of an image. On the other hand, the message can become distorted and consequently be misinterpreted if colour is applied inappropriately.

2.1. Colour Selection

To have proper control over the colours used in creating an image, the user must be able to use an appropriate method for colour selection. *Colour selection* is the process by which the user communicates to the system the colour (or colours) with which a particular object (or set of objects) is to be rendered. For example, specifying the "current" brush colour in a "paint" system is an instance of colour selection. The *colour selection system* is the set of operations that must be used to achieve this communication. One type of colour selection system that is commonly incorporated in many interactive picture creation systems is "palette selection," where the user specifies the desired colour by selecting one of several visible colour swatches using a pick interaction device [Foley82, Tanner83]. Another popular technique is to give the user interactive control over the colour of an area on the display. The user then manipulates different "properties" of colour (e.g. hue, saturation, value) via an interaction device, such as a tablet or mouse, until the desired colour is displayed in the patch.

Colour selection is a task that must be frequently performed in many interactive applications. Hence, if such a system's user interface is not optimized for colour selection tasks, a significant overhead in time and effort will occur, increasing substantially the total work needed to accomplish the task at hand. In addition, a user's tolerance for error (i.e. the discrepancies between desired and actual results) tends to increase with the perceived difficulty of the task. Thus, if the task of colour selection is perceived as being tedious and time consuming, the user will generally be less selective about the colours used in the image creation process. Such compromises degrade the effective functionality of the system since the average quality and complexity of the images created tend to be lower. An easy-to-use colour selection system, on the other hand, encourages the user to take full advantage of the wide range of colours offered, consequently enhancing the usefulness of the system.

2.2. Colour Editing

Throughout history, the human visual system and the perception of colour have been topics of much study. Although a lot has been learned, the behavior of human colour perception is still not completely understood. As a result, using colour to enhance graphical images remains very much an art - one that typically takes years of experience to master [Trucken81]. Though guidelines for the proper use of colour certainly exist, the process of colour selection is often one of trial and error. For reasons described in Section 2.4, the initial colour selected frequently turns out to be unsuitable once an object rendered in the chosen colour is integrated into a complex scene. In many applications, it is therefore desirable to provide the user with a suitable mechanism for altering the colours of objects already rendered.

In reality, the editing of an object's colour can be considered to be part of the colour selection process since, technically, it is not until the user decides on the final colour that a colour for the object has been "selected." In this work, however, the task of colour selection will be considered to be the process of specifying exactly one colour (e.g. the colour that is to be next used to display an object). *Colour editing* is defined as the process of changing the current colour of an object already rendered to a new colour. For example, the editing process could involve:

- selecting the object whose colour is to be changed (e.g. with a pick device);
- specifying the new colour with which the object is to be rendered by using the colour selection system implemented (e.g. by choosing from a colour palette);
- invoking a rendering function that causes the selected object to be redisplayed in the new colour.

On the other hand, the colour selection and editing functions could be merged together to form one combined process in which the colour of the displayed object is manipulated in some continuous manner. This could be done, for instance, by using the movement of an interaction device such as a mouse to increase or decrease the magnitudes of various properties of colour (e.g. hue, brightness, etc.).

In some applications, the colour editing process may be restricted by a number of factors. For instance, in systems that do not record object position information, it is generally difficult to select objects after they have been initially rendered. Thus, it may be necessary to perform editing at the time the object is rendered for the first time. At the very least, the system should allow the user to "undo" the previous rendering operation [Foley82, Singh82]. A new colour could then be selected and the object redisplayed.

A technique that might be preferred is one that would initially render an object in some "temporary" colour and then allow the user to fine tune the object's colour using appropriate feedback. Such an editing function, however, may not be feasible in some picture creation systems where objects are rendered using complex antialiasing or pixel blending algorithms [Tanner83]. On the other hand, some progress has been made in overcoming the difficulties in performing colour editing in such systems. Ken Fishkin [Fishkin84], for example, has developed a number of fill algorithms that take into account the fact that an object has been rendered using antialiasing techniques. It is the belief of the author of this thesis that these algorithms could be extended without too much difficulty to provide a highly interactive colour editing tool for antialiased objects. Such a tool could, for instance, be built by taking advantage of "overlay" pixel planes and multiple colour mapping (features found in such systems as the Adage/Ikonas RDS 3000 [Adage82]) to achieve instantaneous feedback for interactive colour manipulation.

A thorough discussion of colour editing issues is really beyond the scope of this work since it is actually the colour selection component of the editing process that is being studied here. The author feels, however, that it is important to be able to distinguish between these two differing functions. The philosophy of separating task interactions into characteristic components allows proper focus to be placed on isolated manageable problems. This discussion on the topic of colour editing also serves the purpose of further emphasizing the significance of incorporating an effective method of colour selection in the user interface. The proportion of time and effort a user spends initially selecting and subsequently editing the colours used in an image is not trivial. A sophisticated colour editing facility will not be very useful if the system the user selects new colours with is difficult to use. On the other hand, an appropriate user interface for interactive colour selection in conjunction with an easy to use colour editing facility will give the user the flexibility needed to deal with the complexities of effective colour usage.

2.3. Characteristics of an Effective Colour Selection System

In order to design an interface for any situation in which humans are to interact with a computer system, it is first necessary to to determine what properties characterize the interface as being "effective." These characteristics can then be used as a yardstick for subsequent evaluation. The properties that can be considered critical in the design of an effective user interface for interactive colour selection are listed below [Tanner83].

- The colour selection system should provide enough functionality to ensure that the user can select all the colours that are needed in a particular application and at the desired level of precision.
- The user should be able to select colours as quickly as possible.
- Colours should be selected with as little cognitive effort as possible and with a minimum of physical movement.
- The method of selecting colour should be easy to learn by users who are not familiar with the system and, in many cases, have little technical knowledge of computer systems.

These properties are the most significant in terms of user interface design and are the primary criteria upon which the study described here is based. Some characteristics of secondary importance that enhance the user interface in terms of the properties mentioned above are also given.

- The colour selection system should be well integrated into the structure of the overall user interface. That is, the task of selecting colours should be consistent with other interactions the user performs.
- The interface should be predictable. The operations necessary to achieve the selection of a particular colour should be easily anticipated by the user. One common way of accomplishing this is to present a conceptual model or metaphor that the user can easily comprehend and remember.
- The interface for colour selection should be kept simple. In many applications the user will have to remember numerous other interaction procedures. By maximizing simplicity the user will more easily learn to use the system.
- The system should provide adequate feedback all the time and at all levels of interaction. Inadequate feedback often causes unnecessary frustration.
- The system should appear robust; that is, the system should respond in a sensible manner to abnormal or erroneous user actions. For instance, a colour specified outside the gamut of the display device should be handled adequately.

2.4. Colour Selection Problems

This section analyzes some of the issues that make the task of colour selection difficult. While some of the problems discussed are particular to computer graphics applications, others are more universal and occur as a result of the complexities of the human visual system. The discussion presented here is not meant to be comprehensive; rather, the main objective is to provide the reader with an overview of what the difficulties are and, in so doing, emphasize the complicated nature of designing an effective user interface for interactive colour selection.

One problem that is inherent in many computer graphics systems is the overwhelming multitude of colours made available to the user for image rendering. In a 24 bits per pixel system, for example, there are over 16 million colour "values" that can be displayed at any one time. In addition, if the system uses 10 bit wide colour lookup tables (as in the Adage/Ikonas RDS 3000 system [Adage82]) these 16 million colours can be chosen from over 1 billion available values. It is known, however, that the human visual system can only discriminate in the neighbourhood of 4 million colours (and considerably less than this in the range of a typical colour monitor) [Cowan83b]. Since there are considerably fewer perceivable colours than the graphics device is capable of displaying, the selectable colour gamut can for all intents and purposes be considered continuous. The user faces the task of specifying from this enormous selection one colour for each rendering operation. Doing this in terms of the basic colour coordinates of the display device is clearly not a trivial task, especially for inexperienced users.

To overcome this situation, the number of colour values made available to the user is often reduced to a more manageable level; for example, by requiring the user to select from a standard "palette" of colours [Smith82, Plebon82]. This technique, however, tends to restrict access to the full potential of the system, especially if the user has no control over the contents of the palette. A method that would be preferred in many applications would be one that is easy enough to use but allows colours to be selected at less restrictive levels of colour precision. Such a scheme would be necessary anyway in palette based systems where users are allowed to choose the colours that make up the palette.

It is apparent that in order to construct such a system for selecting from a large number of possible colours, the user must be provided a higher level model of colour with which to work. The problem is that no consensus exists on what this model should be [Cowan83b, Boynton79]. Many of the existing models, or colour order systems, seek to arrange colours in an order that is similar to the order that is intrinsic in the human colour perception process. Due to the general lack of knowledge on the complexities of the human visual system, however, these systems have only succeeded in approximating such a truly perceptual representation. In attempting to design an effective colour selection system it is therefore necessary to incorporate one of these apparently imperfect models. Which will work most effectively for the purpose of colour selection has been an issue of considerable debate and controversy and is one of the major topics of discussion in this work. A detailed description of a number of colour order systems is given in Chapter 4.

Another problem that indirectly affects the colour selection process is the alteration of the appearance of colours when an image is displayed using a medium other than that which it was originally created. Although such discrepancies in colour appearance are often caused by miscalibration of display devices and lack of standardization in image storage formats, problems also occur because the colour gamuts of different display media do not generally coincide. For instance, an RGB video monitor is capable of reproducing more blue colours and fewer bright yellow colours than photographic film. A similar statement can be made regarding the transformation of colours when an image is NTSC encoded for video recording or display on a standard television set. These reproduction problems may cause the transformed image to appear significantly different from the original, possibly destroying some of the aesthetic effect that the creator had intended.

As a consequence, yet another burden is placed on the user of the graphics system since these inconsistencies in colour reproduction may have to be taken into account during the colour selection process if the image is to be ultimately transferred to a different display medium. Although the issue of inter-device colour compatibility is not directly examined in the current study, it is one that may have to be taken into consideration in some applications. At the very least, the colour editing facility should be flexible enough to allow modifications to be easily made if images are to be moved to a different display device. For a more detailed discussion on this topic the reader is referred to [Meyer83, Cowan83a, Fishkin83].

In addition to the technically related problems described above, the colour selection process is further complicated by a number of visual effects that become apparent when individual objects are combined to form a complete image. The basis for these phenomena can be summarized by a rule of thumb stating that the colour of an individual object will be influenced by the colour of its surround. In its simplest interpretation, this principle implies that the perceived brightness of an object's colour will be affected by the intensity levels of the surrounding colours. For example, a colour perceived as "orange" on a completely black background will appear to be "brown" when viewed on a white background. These contrast effects also cause the perception of an object's hue to be affected by the hue of surrounding colours. A dramatic example of this effect is that when a yellow object is viewed on a blue background, it will actually be perceived as orange. In addition, the colours of the surround also have an effect on the perceived size of objects. For instance, a white object on a black background will appear to be larger than a black object of equal dimensions viewed on a white background.

The practical implication of these contrast effects is that the appearance of an individual object may change considerably when the colours of the surrounding environment are altered. However, the presence of these effects in an image does not necessarily mean that the effectiveness of the image will deteriorate as a result. In fact, these effects can be taken advantage of to create a number of useful illusions such as spatial depth cues and warm/cool contrasts. In addition, contrast effects are also instrumental in determining such aesthetic qualities as colour harmony, balance, rhythm, proportion, directionality and motion. It is important to note, however, that even when these effects are anticipated, their magnitude is hard to predict since, in general, the total influence that one object will have on another is dependent on a number of factors including the relative sizes of the objects and the distances separating them. Thus as the complexity of a scene increases, the number of variables that affect the perception of the entire image grows very quickly. This makes the task of determining the "right" colour for an object very often a process of trial and error.

It is apparent that the process of selecting the colour that is ultimately wanted in an image is made more complicated because of these contrast effects. Although the task of colour selection is not directly affected since the resulting problems are really colour editing issues, it is important to note that the more often the user must make changes to the colours used in an image, the more often colour selection must be performed. Hence, the need for an effective colour selection system in many interactive applications becomes quite critical once it is realized how frequently the user must perform colour selection tasks.

2.5. Previous Work

This section gives an overview of previous research efforts in the computer graphics community that are related to the user interface issues of interactive colour selection.

The bulk of the work done so far on the problems associated with colour selection has been primarily focused on the evolution of appropriate higher level colour order models and on the development of computationally efficient algorithms for transforming colour information represented by these systems to device displayable form (and visa versa). [Joblove78a], for instance, describes several such colour order systems that are considered suitable for "user specification of colour" including a cyan-magenta-yellow system for emulating a subtractive colour mixture environment and a variety of systems derived from the perceptual properties of hue, chroma and value originally formalized by [Munsell39]. [Smith78a] also presents two systems of the hue-chroma-value family and develops transformation algorithms for these that are designed for efficient implementation on a computer system. In addition, the Graphics Standards Planning Committee has adopted one of the systems of this type, known as HSL (hue-saturation-lightness), as a standard for colour specification [GSPC79]. A number of transformation algorithms are found in [Fishkin83] including one that converts HSL colour information to RGB (red-green-blue) coordinates. This particular transformation claims to be considerably faster than the original presented in [GSPC79]. In [Meyer80] and [Meyer83], the original Munsell colour notation and the Optical Society of America Uniform Colour Scale are described in terms of their application to a variety of problems in computer graphics. For instance, these colour systems are considered to be well suited for organizing the contents of a colour palette since they approximate perceptual uniformity; that is, numerically equal distances in the representation systems imply perceptually equal differences in colour sensation. A number of the colour representation systems found in these works are described in Chapter 4.

In spite of the abundance of existing high level colour order systems, empirical comparisons of the relative effectiveness of different colour selection systems are quite scarce. One study that has been done in this regard evaluates three colour selection systems: one based on the RGB colour model, one based on HSV (hue-saturation-value) and the other a newly developed colour naming system [Berk82]. The results of this experiment were compiled on the basis of how close the colour of the first "guess" was to the target colour. Using this criteria, the naming scheme was ranked best, the HSV-based system second and the RGB-based system third. The implications of these results for graphical applications that allow the selection of colours using a relatively fast feedback mechanism is clearly limited since the experiment was explicitly designed to investigate colour selection tasks in programming-like situations where the user (i.e. the programmer) does not have access to a communication device (such as a graphics terminal).

The issue of faithfully reproducing colour images on different display media, briefly described in Section 2.4, is discussed in [Meyer83]. In that work and in [Fishkin83], a report on the application of colorimetry principles to this problem is given. In particular, these works discuss using the CIE XYZ colour system as an intermediate representation standard to simplify the inter-media transfer of colour information. [Cowan83b] presents a methodology for calibrating colour RGB monitors in terms of the CIE XYZ system.

The principles of how colours perceptually interact with one another and the implications for computer graphics applications are described in [Trucken81]. Some of these perceptual phenomena and their effect on colour selection issues were discussed in Section 2.4.

3. Methodology For Problem Analysis

The previous chapter defined the problem of colour selection in terms of human-computer interaction and discussed some of the issues involved in designing an effective user interface for this purpose. In this chapter, the basic methodology used in this study to analyze the task of interactive colour selection is described. The initial phase of this study is a psychophysical experiment based on the method of colour matching that was conducted in the Computer Graphics Laboratory at the University of Waterloo.

Briefly, *colour matching* is a task where the subject of the experiment is required to make judgements on whether the colour of two halves of a symmetrical test stimulus are perceptually the same [Boynton79, Wyszecki82]. In the current experiment, the subject was given direct control over the colour of one of the test stimulus halves and manipulated the colour of this subfield until it was the same as the colour of the other half of the test stimulus. Each *session* of the experiment consisted of one subject repeating this task a number of times for several test colours. The colour of the stimulus subfield under the control of the subject was manipulated through the use of a tablet, a graphical input device. One of twelve colour selection systems evaluated in this study was employed for the duration of each session. A more detailed account of the experiment and a description of the colour selection systems evaluated are given in later chapters.

The first section of this chapter discusses at length the objectives that this experiment was designed to achieve. In the second section, the physiological and psychophysical methods of investigation are contrasted. Section three discusses colour matching and why this psychophysical technique is an appropriate basis for the evaluation of colour selection systems.

3.1. Experiment Objectives

The following is a discussion of the main objectives that the experiment was designed to achieve. These goals are listed immediately below and then discussed in detail. The objectives are:

- to empirically evaluate a number of colour selection systems to determine how well each is suited for human-computer interaction;
- to provide a foundation of empirical evidence based on the process of colour matching that will appropriately direct and focus subsequent investigations into the problem of colour selection;
- to provide an environment in which the interactions of subjects performing colour matching tasks can be carefully recorded and subsequently analyzed in detail so that the components of interaction that are important for designing an effective user interface for colour selection

can be isolated;

- to exploit the versatility of computer graphics technology in order to provide opportunities to acquire new insights into the complex behavior of the human visual and colour perception systems.

3.1.1. Comparative Evaluation

It was noted in the previous chapter that there is practically no evidence on the relative effectiveness of different colour selection techniques. As a result of the lack of experimental evidence favouring the use of one method or the other, virtually all current designs of colour selection interfaces tend to be based on intuition. In particular, colour selection systems based on the hue-saturation-lightness family of colour order models seem to be favoured in much of the computer graphics literature [GSPC79, Smith78, Fishkin83]. It is therefore evident that there is an urgent need for empirical data that would indicate whether the use of such systems should be supported or discouraged. This experiment has been designed to provide such evidence.

Specifically, two important elements of the colour selection interface were examined: the colour order system, which permits the user to perform colour manipulations in terms of a "high level" model, and the interaction technique with which the user actually interacts with the graphics system to select colours. A number of colour order models were employed in order to determine which of these models is most suitable for human-computer interaction. In addition to an HSL-based model, several other frequently used colour models, plus a "new" model based on a theory of human colour perception, were evaluated. A number of techniques for interacting with the graphics system were also examined. A detailed description of the colour order systems and interactive techniques that were evaluated is given in the next chapter.

This comparative evaluation of the various selection systems was primarily based on the "characteristics of an effective colour selection system" listed in Section 2.3. Emphasis was placed on the "easy to use" and "easy to learn" characteristics since the criterion of being able to select colours at the desired level of precision was effectively guaranteed for all selection systems (see Section 4.1.3). To facilitate analysis of the learning component, the main focus of the experiment was on inexperienced users. That is, most of the subjects who performed the experiment had had little or no experience with colour mixing techniques or using graphical interaction devices such as tablets. The performance of very experienced users was also examined but in a less formal manner.

It is appropriate, at this point, to qualify what the implications of the comparative aspect of the experiment's results are expected to be. Since it is the nature of experimentation in general to observe substantially simplified phenomena under controlled conditions, the applicability of the results from any experiment, including this one, is inherently limited. Hence, this experiment is not expected to prove (or disprove) outright, for instance, that one system is "better" than all others in all possible interactive situations. It is, however, expected that the experiment will fulfill an immediate need for measures of performance for different colour selection techniques that can be used, in a practical sense, as an aid for user interface design. Given the assumption that the task of colour matching under experimental conditions is appropriately analogous to "real world" colour selection situations (this issue is discussed at length in Section 3.3), it is expected that any statistically significant differences observed in the performance of the colour matching task would have considerable implications for the design of future colour selection interfaces.

3.1.2. A Basis for Future Investigations

Since the experiment described here is the starting point for a study of problems related to colour selection, it has been designed to determine relatively large differences in interactive performance between the selection systems evaluated. Thus, the scope of the experiment is rather wide. For instance, each of the colour order models investigated are representative of a "family" of similar systems. The HSV model evaluated, for example, is but one of many systems all based on concepts similar to the colour properties of hue, chroma and lightness. It was intended that any differences in the performance of the experiment task found to be significant would become the basis of subsequent investigations that, for example, would compare a number of colour representation systems within the same family. A retrospective discussion of possible extensions to the work begun in this investigation is given in Chapter 8.

3.1.3. Determining Critical Components of Interaction

The methodology used to compare how well the different systems are suited for performing colour selection is fairly straightforward. A number of subjects executed colour matching tasks under experimental conditions and the interactions that occurred were recorded. To quantify the relative performance of the various selection systems, a number of performance measures were calculated and each technique was rated according to several criteria.

This type of analysis has been applied in many studies of human-computer interaction and certainly plays an important part in the current investigation. Analysis of this sort, however, is usually most effectively applied at the evaluation stage of interface design where it is already known what the "better" techniques are and it only remains necessary to determine which one of these is "best" [Card83]. More precisely, a comparative evaluation only allows one to make statements such as:

> According to criterion X, method A is N units more effective for performing the investigated task than method B.

Although such conclusions can be made for each method included in the evaluation, very little can be deduced about other possible methods. Since our knowledge of the human colour perception process is far from complete, it is not really sufficient to base an extensive study such as this one solely on a comparison of existing colour selection techniques since it is quite likely that an "optimal" system for colour selection has yet to be developed.

It is primarily for this reason that one of the important goals of the current investigation is to determine which components of interaction are significant for the execution of colour selection tasks. Such knowledge would allow for a fundamental understanding of the colour selection process and greatly increase the chances for the development of a truly "optimal" selection system. A purely comparative analysis, by itself, does not explain *why* method A is better than method B. If the reasons for method A's relative success can be discovered, then it is quite possible that method A can be further enhanced.

At this initial stage in the study, the effort to determine which components of interaction are significant for selecting colours has been carried out by using a variety of strategies to analyze the data collected from the colour matching experiment. For instance, in one approach the "path" of interactions for individual trials (i.e. one colour match) was examined to discover instances in which a subject appeared to have particular difficulty (or ease) in performing the required task. If such a phenomena appeared to reoccur with some regularity in other trials, then a focused statistical analysis was performed across the remainder of the data to objectively determine the degree to which the phenomena affected the entire group (or selected subgroups) of subjects. Patterns in task performance were also brought to light by applying measures that intuitively seemed to be good indicators for the performance of subjects at different stages of the colour matching process. For example, measuring the time it took subjects to reach different thresholds of colour difference provided useful insight in this regard. A detailed discussion of the analysis performed on the colour matching data is presented in Chapters 6 and 7.

3.1.4. Discovering New Insights

The goal of one area of colour vision research is to determine a colour model that "orders" colours (in a mathematical sense) in a manner that emulates the behavior of the colour perception system [Cowan83b]. Although there is considerable debate on whether one model can in fact achieve this, there is a school of thought which suggests that the model that would be considered "perfect" for the purposes of human-computer interaction would be a model that mimics the behavior of human colour perception in this way. The converse of this implies that if one could indeed develop a colour selection system that was found empirically to be very easy to learn by even the most naive of users, then one could argue that the model employed for human-computer interaction may in fact be closely approximating the internals of the human visual system. To be able to make such a statement with reasonable confidence would certainly be a significant breakthrough in our understanding of how the human perception system processes colour information.

An attempt has been made in the current study to contribute to the enormous task of determining exactly how the human visual system works. A colour order model based on the "opponent colours" theory of human colour perception has been incorporated in several of the colour selection systems included for evaluation in the experiment. The opponent colours theory is believed by most members of the colour vision community to be a good approximation of how colour information is encoded at an intermediate neural level before being sent to the perceptual centers of the brain [Boynton79, Cowan83b]. (This model is described in some detail in Section

4.2.5 of this work). It is in the implementation of this model, which would normally require the design and construction of an expensive electro-mechanical prototype, that the inherent versatility of computer software has really been effectively "exploited." An efficient and practical realization of the model, which may otherwise have remained a theoretical entity, has been achieved in a relatively simple manner. Furthermore, this kind of implementation of the model allows any subsequent enhancements to be made by merely modifying the software.

In some sense, then, the quest for a theoretically sound model that accurately accounts for all observed colour perception phenomena and the search for the "perfect" system with which to perform interactive colour selection are the same. By harnessing computer graphics technology to implement some of the existing models of human colour perception, there is, in fact, a unique opportunity for learning more about the perception of colour. It is not expected, however, that great strides towards obtaining any dramatic insights in this regard will be made in the short term. In the long run, though, it is quite possible that some form of empirical evidence can be provided to support (or contradict) some of these theories.

3.2. Advantages of a Psychophysical Approach

Two standard strategies are commonly used to investigate the human visual system: the physiological approach and the psychophysical approach [Boynton79, Wyszecki82, Cowan83b].

The physiological method is a "bottom up" approach that concentrates on the physical components of the visual system and on the relationships between these components. A physiological investigation of the visual system typically begins by studying the responses of individual cells to certain visual stimuli. By then analyzing the interactions that occur with surrounding cells in a hierarchical fashion, the physiological response to a particular stimulus can be traced up the visual pathways towards the perceptual centers of the brain.

The major advantage of the physiological approach is that it is a very systematic technique that, in theory, can lead to a complete understanding of the visual system. In practice, this technique has had most of its success in the analysis of relatively homogeneous organisms that are characteristically composed of a single element type repeated many times in some regular manner. The physiological strategy has, for example, lead to important insights into the functioning of the retina and the lateral geniculate nucleus, an intermediate relay mechanism that transmits optical data received from the eye to the visual components of the brain.

The networks of cell elements that must be analyzed in order to achieve a practical understanding of the visual system are, however, much too complicated to even consider a physiological study for the purposes of this investigation into human-computer interaction issues. Critical parts of the visual pathway are made up of far too many components of different types to allow for a meaningful subdivision of the system into functional modules. Such a breakdown into subunits is vital to developing even a primitive understanding of the visual system's behavior. The microscopic size of the components that must be examined also makes the complex connectivity patterns very difficult to trace. Furthermore, many of the analytical techniques that are used in a physiologically based investigation cannot be directly performed on human beings. Hence, it is often necessary to conduct experiments using animals and subsequently infer conclusive results about the human system.

The psychophysical approach treats the visual system as an unknown entity and deals only with the relationships between the input and output of this "black box." Typically, the visual environment of an observer (the input) is manipulated in some fashion while the perceptual sensations of the observer (the output) are somehow measured and recorded. Behavioral models are then developed that seek to explain the characteristics of the observed input-output relationships.

Traditionally, the ultimate objective of psychophysical analysis of the visual system has been to determine principles of behavior that are relatively general in nature and can be used to explain a number of observed phenomena. One of the major advantages of the psychophysical approach, however, is that it is in fact very easy to study the responses of the visual system in a way that is relevant to a particular application. It is primarily for this reason that the psychophysical method was adopted for the current study. Colour selection is a relatively specific task and being able to measure and compare the performance of subjects using different colour selection techniques without having to explain observed behavior in terms of "neurons and photoreceptors" is clearly advantageous. The psychophysical approach also has the important advantage that virtually all analysis can be performed directly on humans. Any results determined to be significant are therefore much more applicable to human-computer interaction problems.

One potential disadvantage of the psychophysical method is that this system of investigation is not very well defined. The researcher seeking to determine new insights about the visual system faces the very difficult task of choosing the "right" combination of visual stimuli and response restrictions for empirical analysis. Occasionally, the analyst will uncover a fundamental relationship that occurs in many other visual circumstances. Often, however, the investigation will merely produce a complicated visual function that is specific to the inputs and outputs studied. The input and output parameters chosen for the current investigation are described in the next section.

3.3. Colour Matching

One of the simplest and most fundamental psychophysical techniques for investigating the human colour perception system is the colour matching experiment. As in most psychophysical investigations, the inputs and outputs of a colour matching experiment are kept very basic. The visual stimulus consists of a symmetrical display field that is subdivided into two identical subfields. The colour of one subfield is fixed while the other is variable. The subject's perception of the visual environment while fixating on this stimulus is measured by replying "yes" or "no" to the question: "Do the two colours look the same?" This response is the simple basis of the experiment upon which models of visual behavior must be built.

Colour matching experiments have been successfully used to determine a wide range of valuable facts about the human colour perception system. For instance, the relative spectral sensitivities of several components of the visual system, including the rods and cones (the physiological organisms located in the retina of the eye through which light energy is initially detected), have been determined through such experiments. In fact, much of the existing knowledge about colour vision has been derived from the results of colour matching experiments.

Although the basic operation employed in most colour matching studies is fundamentally the same, there are frequently variations in some of the procedural details of the experiment. Minor differences in the shape and orientation of the test stimulus, the conditions in which the experiment is conducted (e.g. ambient lighting) and the actual task subjects are required to perform are common. In a variation of the colour matching task which is employed in the current experiment, subjects are given direct control over the colour of the variable subfield. Each match then becomes a *sequence* of discrete decisions: "No match yet, no match, no match, ... yes the colours are matched!" More precisely, since it is the subject that must decide how to change the colour of the stimulus, each match can be thought of as a sequence of "colour points:" colour t0, colour t1, colour t2, ... colour tN (where "colour tI" is the colour of the variable subfield at time I).

Historically, most colour matching investigations of this type have utilized only the final element in this sequence of responses. That is, the colours of the two display subfields (or the difference between them) at the end of each match are the only pieces of information of the entire matching process that are recorded. The primary reason that the remainder of this information has been ignored is that most previous investigations have simply not been interested in data other than the final stimulus colours. Another important factor may be that until very recently recording the state of the test stimulus prior to the end of the match was extremely difficult. However, since the current experiment was implemented using a computer system, this information was obtained very easily.

As it turns out, colour manipulations of the variable subfield that occur before the match is completed are of critical significance in the current study on interactive colour selection. In fact, it is only by examining these interactions that one can achieve the goal of determining which components of the colour selection process are important for user interface design. For example, this data allows for the easy calculation of the length of time it takes subjects to reach various "thresholds of colour difference;" that is, how long it takes to manipulate the colour of the variable subfield so that the difference in colour between the two subfields of the test stimulus is less than a particular fixed value. (This measure of subject performance is discussed in greater detail in Chapters 6 and 7). This emphasis on the colour manipulations that lead up to the completion of a match rather than on the final results themselves is one of the significant differences between the current experiment and traditional colour matching.

Another point of departure is that in the current investigation it is not taken for granted that the final colours of a match are perceptually the same. Since it is expected that different colour selection systems will have varying degrees of "ease of use", it may very well be that subjects using different systems will, on average, match the two colours to different levels of precision. Hence, one of the primary performance measures used in the analysis of experiment's results is the perceptual difference between the two test colours at the end of a trial. Since there is a strong possibility that there is a significant tradeoff between the time taken to complete a match and how close the colours are matched, this criteria must be taken into account when comparing colour selection systems.

The form of colour matching used in the current experiment offers a number of advantages that make this task suitable for this initial investigation into interactive colour selection.

First, the historical success of the colour matching method has clearly demonstrated the usefulness of this technique as a fundamental tool for investigating colour vision phenomena. Furthermore, the procedural "rules" for conducting this type of experiment are very well defined.

Secondly, a colour matching operation in which the subject is given control of one of the test colours is appropriately similar to "real world" colour selection situations. The task performed by subjects in the experiment is, in fact, a restricted form of colour selection. The primary difference between the version of colour selection performed in the experiment and that performed in most computer graphics applications is that in colour matching the colour that is to be selected is *given*. In more realistic circumstances it is the user who decides what colour is to be selected. Furthermore, it is often the case that the user has only a vague notion (or perhaps no idea at all) of what colour should be used in the rendering process. Hence, experimentation is often necessary to determine the final colour. However, as was previously discussed in Section 2.2, how easy it is to perform such experimentation is a "colour editing" issue as far as this investigation is concerned and is therefore not directly examined in this study. Another dissimilarity with more general colour selection tasks is that in the experiment subjects are requested to match colours "as close as possible." Not all applications require colours to be selected to this level of precision. Nonetheless, a meaningful comparison of performance can be obtained for lower precision requirements by measuring the time it takes subjects to reach different closeness thresholds. Such measurements would, for example, be indicative of how long it would take users to select a colour that is in the "general neighborhood" of the target colour. In spite of the discrepancies described above, it is felt that the interactions executed by subjects performing the colour matching task are consistent enough with those interactions occurring in "realistic" situations that the objectives of this investigation can be achieved.

The third advantage in conducting a colour matching experiment is that task performance can be easily measured in an objective manner. If it were the case that subjects were permitted to select arbitrary colours, then each subject would in effect be performing a unique task making an empirical comparison of selection systems extremely difficult. Requiring subjects to select the same set of colours, however, makes the experiment task uniform (with the exception of the selection system used) across subjects and creates a standard against which meaningful performance criteria can be measured.

4. Colour Selection Systems

There are conceivably many ways for a user to select colours in a computer graphics system. Several techniques that are commonly used for performing colour selection, including the selection systems that were evaluated in the colour matching experiment, are described in this chapter. In the first section, a taxonomy of commonly used colour selection techniques is presented. The second section describes the concept of a colour order system and discusses each of the systems models evaluated in the experiment. In the third section, two techniques for manipulating the colour attribute parameters of a given colour order system are described. Both these interaction techniques and the five colour order models described in the second section form the basis of each of the colour selection systems evaluated in the experiment. A precise description of these selection systems is given in the fourth section.

4.1. Methods of Colour Selection

Most of the methods commonly used for performing colour selection tasks can be classified into three distinct categories: colour naming, palette selection, and numerical specification. In this section, each of these categories is defined and the advantages and pitfalls of each discussed.

4.1.1. Colour Naming

In a colour selection system that is based on colour naming, the user describes the desired colour using words, or phrases, of the English language [Cowan83b]. Such naming schemes are often based on commonly used terminology. Many colour naming schemes have been proposed and several have become standards for use in a variety of different applications [Kelly55, Maerz50, Berk82].

The effectiveness of a colour selection system based on colour naming is largely dependent on the number of colours that must be discriminated in a given task. It is generally thought to be convenient to reference colours through colour naming when the number of colours that must be distinguished is on the order of 10. For instance, a user who selects a colour labeled "red" out of a total set of 10 colours is probably willing to accept a wide range of colours that actually meet the expectation of what "red" should be, since the other colours in the set are likely to be "green," "yellow," "orange," etc. Selection systems based on colour naming are particularly useful in situations where the user has no immediate access to feedback that indicates which colour has actually been selected (e.g. when using systems such as the Apple Macintosh that provide only a monochromatic display but allow colours to be selected for multicolour hardcopy devices [Apple84]).

As the set of selectable colours grows, however, the difficulty in using such naming schemes to distinguish between colours becomes increasingly difficult. In fact, there seems to be a fundamental limit in the human cognition system restricting the number of names that can be used, even by an expert user with years of practice. This limit appears to be about 1000 names [Cowan83b]. Thus it is apparent that for tasks that require users to select from and distinguish between a much larger set of colours, there is need for an alternative method of colour selection.

4.1.2. Palette Selection

The colour selection technique called *palette selection* is defined here as one in which the user achieves the specification of a given colour by pointing to a displayed instance of the desired colour and selecting it using a pick interaction device such as tablet or mouse [Tanner83].

Numerous variations of this technique are in common use in interactive computer graphics applications. One technique that is frequently used is to display a matrix of colour swatches along the perimeter of the screen area [Plebon82]. The user then simply "picks" the swatch having the desired colour, which then becomes the current rendering colour. Often this palette of colours is not shown during drawing operations and "pops up" under user control [Smith82]. Many applications also allow the colours contained in the palette to be defined by the user, in which case an alternative colour specification scheme is used. Some systems permit users to pick any currently displayed colour. Thus, any area on the display screen, including the image itself, can be used to select colours. Another technique, that to the author's knowledge has never been implemented but has some intuitive appeal, is to provide a hierarchy of palettes for selecting colours. Such a system would attempt to get around the problem of being able to display only a small number of selectable colours by displaying a successively more "refined" palette for each colour selection made by the user. For example, after the selection of an "orange" colour from the main palette, the system would display a subpalette that consists of colours that perceptually appear "in the neighbourhood" of the previously selected colour (e.g. various shades and tints of orange). One major disadvantage of this technique, however, is that, for many applications, the hierarchy would have to be accessed 3 or 4 levels deep for each colour selection In addition, when the user discovers that the desired colour is not on the currently displayed palette then the additional overhead of ascending to the previous levels in the hierarchy must be incurred.

For many tasks, the palette selection technique is the most convenient method for performing colour selection since the user need merely point at the target colour and invoke the selection mechanism. In addition, this type of colour selection system is often used in order to provide convenient access to a manageable subset of frequently used colours. Thus, it appears that palette selection schemes are most useful in situations where the number of colours used is relatively small and the colours included in the palette can be conveniently preselected. In many applications, however, it is necessary to have easy access to very large number of colours. It is then necessary to have a specification system that allows for the convenient selection of colours at a high level of precision.

4.1.3. Numerical Specification

The user who selects colours using a system based on *numerical specification* is presented a mathematical model, or *colour order system*, with which colours can be selected to any desired precision [Cowan83b]. The topology of visible light, or colour space, necessarily implies that such a model must have at least 3 dimensions, or degrees of freedom, if all colours are to be specifiable. Thus to select a colour using such a numerical technique, the user must explicitly or implicitly specify a point in 3 dimensional space; that is, one coordinate value for each degree of freedom. The main disadvantage of this technique is that it is very difficult for a user to remember the exact mapping from the numerical values to perceptual sensations of colour; that is, it is relatively hard to determine which numbers produce which colours. Thus it is generally not desirable to require the user to explicitly specify these numeric values, especially for tasks where it is necessary to select colours at a relatively high level of precision.

The primary goal of a colour order system is, however, to *order* the set of visible colours in a way that is conceptually meaningful in terms of the way human beings perceive colour. That is, colours that are judged to be very close together by human observers should lie closely together in the colour model, while colours that appear to be very different lie far apart. These models define an ordering of colour space such that a gradual progression in any particular "direction" within the defined colour order system will correspond to a smooth (but not necessarily linear) transition in the perceived colour. This implies that a user should be able to predict in which direction a desired colour lies and consequently narrow the "search" for the target colour deterministically.

It is apparent that the effectiveness of such a "trial and error" strategy is dependent on a number of factors, one of the most important of which is the duration of the period between the time a colour is specified and the time the specified colour is actually displayed. Up to a certain threshold, the faster this feedback mechanism the faster a user will be able to conduct the trial and error process. One technique that uses feedback that is effectively instantaneous and does not require users to explicitly specify numeric values uses a continuously variable interaction device (such as a potentiometer, slider, tablet, mouse, etc.) to specify the current position in colour space. This method permits users to manipulate colour in a smooth and continuous manner and allows the trial and error process of selecting a desired colour to be performed very efficiently.

The primary advantage of using a numerical specification technique is that users can specify colours at any desired precision and can therefore select any colour the display device is capable of producing. Given that the user can implicitly specify the coordinate values via an interaction device, the major difficulty in designing a user interface that implements this technique is deciding which of the many existing colour order models should be used as the basis of the selection system. In fact, it is by no means clear whether the colour model used actually makes a difference in how easy it is for users to select colours. Determining whether or not users can perform colour selection tasks more easily when using different colour order systems and, if so, determining which colour models are better suited for incorporating in a human-computer interface are two issues which the current study is intended to help resolve.

4.2. Colour Order Systems

In the current experiment, five colour order systems are evaluated in an attempt to determine which, if any, are suitable for performing colour selection tasks in the manner described above. Each of these is now described in turn. A detailed account of how these colour models were implemented is given in Appendix A.

4.2.1. RGB

The RGB colour order system is a coordinate space based on the properties of additive colour mixing [Cowan83b, Fishkin83, Smith78, Berk82]. According to these principles, several "primary" colours can be mixed together in different proportions to create new colours. Experimentally it has been shown that three properly chosen primaries can be used to produce most colours that are perceivable by the human visual system. Roughly speaking, a suitable set of primaries consists of one primary from the red part of the spectrum, one from the green part and one from the blue part. Hence in the RGB colour order system, the "R," "G" and "B" coordinates specify the relative proportions of the red, green and blue primaries that are mixed to form the resultant colour. This principle is in fact the mechanism used in many display devices, such as CRT monitors, to produce colour. In the colour matching experiment, the three primary colours used to implement the RGB colour order system are exactly the colours of the red, green and blue phosphors used in the colour monitor that displayed the test stimulus.

One of the characteristic problems associated with displaying images through a medium such as a video monitor is that the relationship between the driving voltages of the electron guns and the actual amount of light emitted from the phosphors is not linear. The implementation of the RGB colour model used in this experiment has been corrected for such non-linearity using the "gamma correction" procedures outlined in [Cowan83a]. Thus the coordinates of the RGB system used by subjects to perform colour matching are linear with respect to the excitation levels of the three phosphor types.

The use of RGB order systems that are based directly on phosphor excitations is very common in computer graphics applications since this colour model is inherently available on any system that uses a CRT type colour monitor as an output device. Even with the additional complexity of providing gamma correction, colour selection systems based on this model are very easy to implement. In addition, many system implementors are partial to the RGB model because many users are already familiar with it. The RGB system is one of the standard colour systems adopted by the ACM/Siggraph Graphics Standards Planning Committee for the CORE graphics system [GSPC79].

On the other hand, the RGB colour model makes little attempt to take into account any of the perceptual characteristics of the human visual system and consequently is commonly believed to be counterintuitive for user interaction purposes (especially for naive users). For example, linear changes in RGB coordinate values generally do not imply linear changes in *perceived* colour.

4.2.2. YIQ

The YIQ colour system is the colour scheme that was initially proposed by the National Television Standards Committee (NTSC) of the U.S.A. and was subsequently adopted as the standard for transmitting colour information for broadcast television in North America and Japan [Hunt75, Meyer83, Smith78]. The primary objective in the design of this system was to maximize the perceptual resolution of the encoded colour information using a minimum of signal bandwidth.

Since it had been shown experimentally that the human perceptual system can more easily detect differences in luminance (i.e. the amount of light energy emitted) in a given visual pattern than differences in chrominance (i.e. those properties of colour not related to luminance; for example, hue and saturation), the NTSC recommended that a relatively large portion of the available signal bandwidth be devoted to representing luminance information. This implied that the luminance component should be separated from the chrominance information. The Y axis of the YIQ system was therefore defined in order to approximate the luminance property. The separation of luminance also facilitated compatibility with the "black and white" transmission system that was being used at the time. A monochromatic television receiver had only to ignore the chrominance information (i.e. the I and Q axes) to produce images that were equivalent to those produced using the older system.

Luminance is commonly believed to be a separate degree of freedom at higher levels in the colour perception process. Hence the separation of the luminance component in the YIQ colour system is an attempt to isolate a psychophysical property of the human visual system. Most existing colour order systems (in fact, 4 of the 5 evaluated in this study) provide a separate degree of freedom for luminance or a similar perceptual property.

The I and Q axes of this colour order system were also defined according to the criteria of maximizing the information content of the encoded colour data. Roughly speaking, the I axis represents chrominance information in the blue-green (i.e. cyan) to orange direction, while the Q axis represents colour in the yellow-green to purple direction. It was found experimentally that this encoding scheme was best in terms of "blue saving." Blue saving is a technique that takes advantage of the fact that the human visual system is much less capable of discriminating blue colours than colours of other hues. By separating out the blue component, the available signal bandwidth can be used more efficiently. The I and Q coordinates defined by the NTSC turned out to be optimal in this respect for the colour gamuts of television display devices.

In theory, the YIQ colour system is considered to be "standard" in that the same results should be produced on different colour display devices. This standard is based on recommended chromaticity values for the the underlying red, green and blue primaries. In other words, the chromaticity coordinates of the phosphors used in all display devices are assumed to be the same. In practice, however, manufacturers of display devices frequently deviate from these recommended values. As a consequence, the properties actually controlled by the user of a YIQ-based colour selection system are dependent on the colour monitor used and will tend to vary between display devices.

In terms of usability, the YIQ system has been criticized in the computer graphics literature as not being well suited for the purposes of human-computer interaction. Berk et al. point out, for instance: "The system is not attractive at all from a user's point of view - it is certainly not intuitively clear [Berk82]." The YIQ system is included for evaluation in this experiment primarily because of its frequent use in television related industries. In addition, this system is familiar to anyone who has attempted to adjust their home colour television set.

The YIQ colour order system, as implemented in this experiment, is derived from the gamma corrected RGB system using a simple linear transformation. The exact formula used to do this can be found in Appendix A.

4.2.3. HSV

The HSV colour order system [Smith78, Joblove78, Meyer83, Fishkin83] approximates the perceptual properties of "hue," "chroma" and "value" originally proposed in [Munsell39]. Roughly speaking, ***hue*** can be defined as that particular property of a colour that "enables it to be assigned a position in the spectrum [Guralnik80]". Red, yellow and cyan are different hues. Red and pink, on the other hand, are of the same hue but of differing ***saturation***. Saturation is the attribute that quantifies the distance of a colour from the achromatic colours (i.e. white or gray). ***Value*** as defined by Munsell, is very similar to the property of luminance (see Section 4.2.3). However in the HSV system used in the colour matching experiment, the value attribute more closely resembles a measure of a colour's distance from black. This property is related to luminance but is dissimilar in that all pure hues are considered to be of equal value but are not necessarily of equal luminance. Pure yellow, for instance, has approximately twice the luminance of pure red when displayed on a video monitor.

The HSV system used in this investigation is a direct implementation of the algorithm presented in [Smith78] and is reproduced in Appendix A of this thesis. The HSV system is based on a transformation that maps HSV coordinates to the (gamma corrected) RGB system. Thus, like the RGB system, HSV is also defined with respect to the natural coordinates of the display monitor and is therefore also dependent upon the characteristics of this device.

Colour systems of the hue-saturation type are believed (especially within the computer graphics community) to be well suited for performing interactive colour selection tasks. In particular, Smith claims: "in some situations, such as colour mixing, these models are more intuitively satisfying or more convenient to an artist than the colour cube [i.e. the RGB system]". A colour system of this type was selected by the ACM/Siggraph Graphics Standards Planning Committee as one of the standard colour schemes for the CORE graphics system [GSPC79]. As a result, the use of these colour systems in computer graphics applications is very common. The HSV colour space (as presented in Smith78) is often used instead of other colour systems of the hue-chroma-value "family" because it is computationally more efficient than others.

4.2.4. CIELAB

The CIELAB colour order system was adopted as an international colorimetry standard by the CIE (Commission Internationale de l'Eclairage) in 1976 [CIE78, Wyszecki82]. This system was chosen to represent the CIE family of standard colour systems in this study primarily because it is the closest existing approximation to a perceptually uniform colour system. A colour order system that is perfectly uniform in terms of human colour perception is one where differences in colour that are perceived by humans to be equivalent are represented by equal Euclidean distances in the colour system. In terms of human-computer interaction, this implies that there is a linear relationship between the physical movement of an interaction device and the perceived change in colour. It is felt that this characteristic may very well be significant in designing an effective interface for interactive colour selection.

The CIELAB colour system is based on a non-linear transformation of CIE XYZ tristimulus values. (XYZ tristimulus values are used universally as a standard method of representing and communicating colour information). According to its formal definition, the transformation from the CIELAB system to XYZ tristimulus values is done relative to a reference white colour. For simplicity, this reference has been chosen to be "equal energy white;" the point on the CIE chromaticity diagram having coordinates (0.33, 0.33). The exact transformations used to map CIELAB coordinates to XYZ tristimulus values and subsequently to the RGB system are given in Appendix A.

Approximate perceptual uniformity is a property that makes the CIELAB system intuitively appealing for performing colour selection tasks. In addition, CIELAB is not dependent on the properties of display monitor actually used, since this system is based on standard tristimulus values. This means that the behavior of the colour selection system will not vary from monitor to monitor. Another advantage of this system is that, since it is a colorimetry standard, it is already in common use in many non-computer graphics applications.

4.2.5. Opponent Colours

The Opponent Colours model is a theory of human colour perception that attempts to explain how colour information is processed between the instant light energy is first detected in the eye and the moment it is sensed as a colour in the perceptual centres of the brain [Boynton79, Cowan83b, Goetz82]. This model has now been accepted by most members of the colour vision community as a valid explanation of how the intermediate processes of the colour perception system behave.

The essence of the Opponent Colours model is shown graphically in Figure 4.1. According to the model, the perception of colour occurs in two principle stages. Initially, light is absorbed by cones in the retina of the eye. There are three fundamental cone types found in the retina: an R type which is most sensitive to long wavelength light and can be roughly equated to "red;" a G, or "green," type that is most sensitive to medium wavelength light; and a B, or "blue," type that is most sensitive to short wavelength light. The colour signals originating from the retina are then transmitted via three independent neural pathways, or *channels* (one for each of the red, green and blue components), to the second processing stage. Here the signals are "summed" and "differenced" (i.e. in a behavioral sense, not a strictly mathematical one) to create three new neural signals. The signals from the R and G cone types are added to form an achromatic, or luminance, signal. The difference of these two signals is also taken to form a red-green opponent signal. A yellow-blue opponent signal is formed by differencing the signal of the B cones with the output of the achromatic channel. The three outputs from this process are transmitted up the hierarchy of the visual system and eventually end up in the brain, resulting in a sensation of colour.

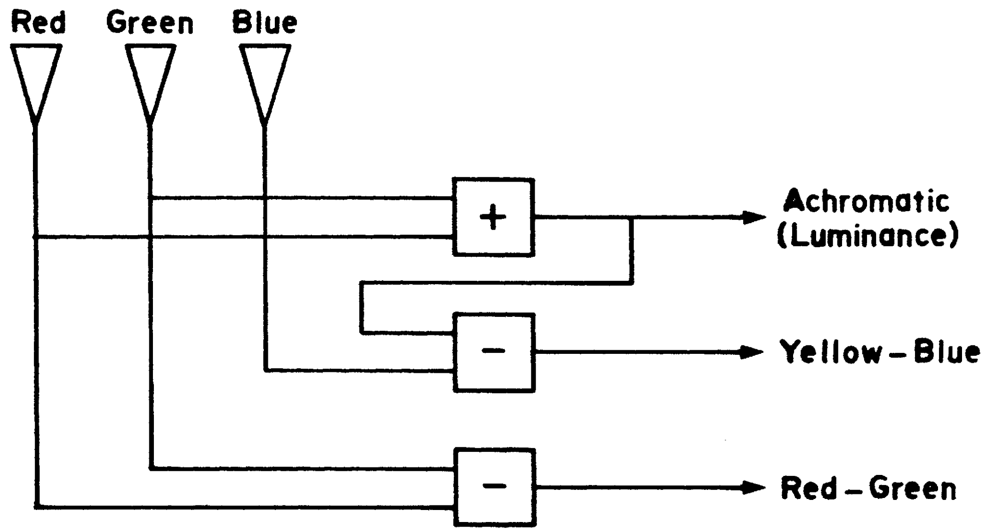

Figure 4.1. The Opponent Colours model.

The Opponent Colours model helps to explain a number of visual phenomena such as why the colours red and green are never seen in the same place at the same time. Variations of this model have been developed to account for other perceptual effects. In fact, one of the major goals of current colour psychophysics research is fine tuning the Opponent Colours theory.

A colour order system that is derived from the Opponent Colours model was implemented for evaluation in the current study. The implemented colour system is based on a transformation of standard CIE XYZ tristimulus values (see Appendix A for a description of how XYZ tristimulus values are transformed into device coordinates). The luminance signal, or axis, is set equal to the Y tristimulus value. The derivation of the chromatic axes is illustrated in Figure 4.2. using the standard CIE chromaticity diagram. The chromatic axes are based on four vectors radiating outward from the white reference point through four points located near the perimeter of the colour solid. These points have been chosen in order to define the hues of "red," "green," "blue," and "yellow." All but the yellow reference point are defined by the phosphor chromaticity coordinates of the display monitor used in the experiment. The yellow point is chosen to have a hue equivalent to light having a wavelength of 580 nm, the hue that is normally equated to "unique yellow" [Boynton79]. The angles between the four vectors are subsequently transformed to produce a right-angled coordinate system as shown in Figure 4.2. The exact algorithm used to implement this system is given in Appendix A.

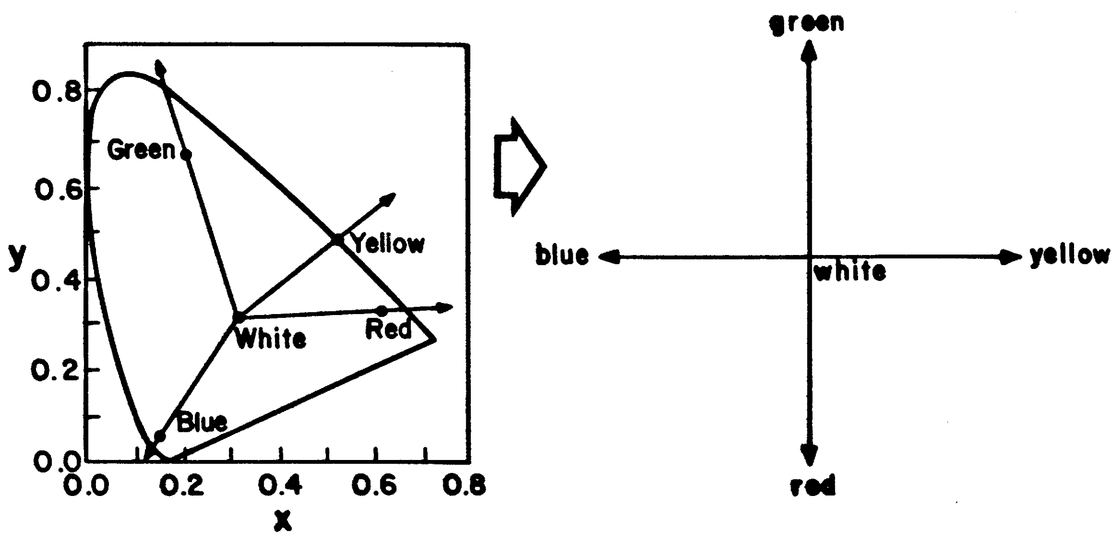

Figure 4.2. The derivation of the Opponent Colours colour order system.

The Opponent Colours system is appealing for a colour selection user interface because of its theoretical basis. However, the model that the system is derived from is not, as yet, fully developed. Furthermore, its current implementation requires a significant amount of computational power (e.g. hardware to assist in the evaluation of floating point calculations) to make a real-time interactive environment feasible. Consequently, a simplification of the algorithm would be desirable if the Opponent Colours system is indeed found to be well suited for colour selection tasks.

4.3. Interaction Techniques

A colour order system, by itself, is not enough to provide a complete human-computer interface. In order for interaction to actually occur, a user must have a means by which the three degrees of freedom of a given colour system can be controlled. In computer graphics this usually implies the use of some sort of interaction device such as a tablet or mouse. In designing an interface, however, another issue arises. Will the user want to alter only one degree of freedom at a time, or will it be useful to be able to manipulate two, or even three, degrees of freedom simultaneously? It may even turn out that the solution to this problem is dependent on the colour order system used.

Two interaction techniques were evaluated in the current study. Both these schemes are implemented using a tablet and four-button puck. The tablet has been shown to be an effective way of emulating many types of interaction control devices [Evans81]. It is largely due to this versatility that a tablet was chosen to interface the interaction techniques described in this section.

Relative changes of puck position are used to compute changes in colour. This emulation of a mouse device is used instead of the normal position sensitive method primarily to facilitate finer control over colour modifications. The position sensitive technique is limited in the sense that the entire range of a controlled parameter must be mapped onto the range of coordinates provided by the physical device. In the case of users performing colour selection tasks, this implies that small puck movements will cause relatively large perceivable changes in colour. This is clearly unsatisfactory for the purposes of this experiment, since subjects are required to match colours as closely as possible. The motion sensitive method permits an arbitrarily large range of "virtual" coordinates to be defined, consequently making more suitable hand motion - colour change ratios possible. For all the colour selection systems included in the current evaluation, each degree of freedom was appropriately scaled so that the entire range of each degree of freedom projected onto an area approximately twice the size of the 32 centimetre square tablet.

4.3.1. Single Attribute Control

The first of the interaction techniques evaluated in the experiment uses an independent control mechanism for each of the three degrees of freedom of a given colour model. That is, each colour property is controlled by a separate button on the puck. Only the horizontal axis of the tablet is used to determine changes in colour; vertical movements of the puck are ignored. Furthermore, all puck movements are disregarded if more than one puck button is pressed simultaneously. Thus only one degree of freedom can be altered at once.

An implementation of this technique for the HSV system is shown in Figure 4.3a. If, for instance, the topmost button is held down and the puck moved to the right, the "value" of the displayed colour will increase. Conversely, value will decrease if the puck is moved to the left while the same button is held down. Hue and saturation are controlled in a similar manner by the left and right buttons, respectively.

4.3.2. Multiple Attribute Control

The second technique utilizes a control mechanism that permits the simultaneous control of two degrees of freedom. In this scheme, two properties of colour are controlled by one puck button, while another button is used to control the third. Simultaneous control of two colour properties is achieved by using both the horizontal and vertical components of puck movements to determine the change of the two attributes. Again, only one button may be pushed at a time; all puck movements are ignored when more than one button is held down.

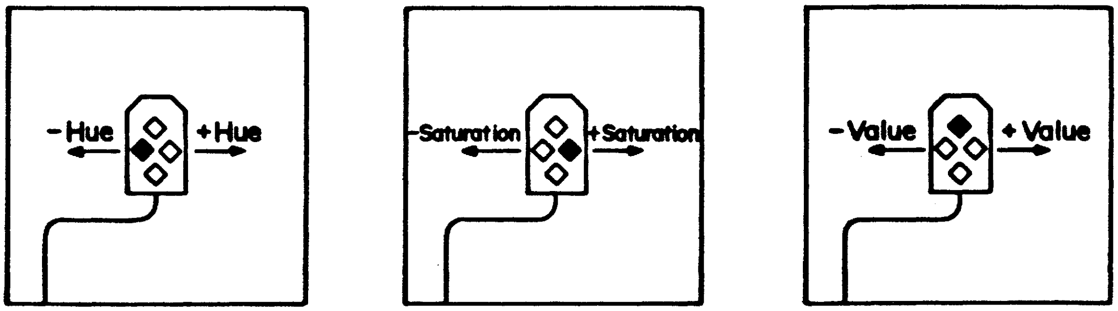

Figure 4.3a. The HSV colour order system interfaced via the single attribute control interaction technique.

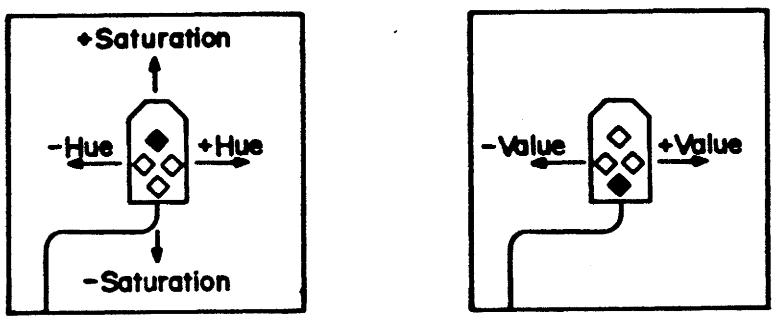

Figure 4.3b. The HSV colour order system interfaced via the multiple attribute control interaction technique.

This technique is illustrated, using the HSV colour space as an example, in Figure 4.3b. When the top button is depressed, movements of the puck in the horizontal direction determine changes in hue while movements in the vertical direction are used to compute changes in saturation. If the puck's movement is precisely horizontal, only the displayed colour's hue is changed. Likewise, only the saturation property is altered for puck movements that are exactly vertical. If the puck is moved up and to the right at an angle of 30 degrees with respect to the horizontal axis, the hue property is "increased" at a rate twice that of saturation. The value property, in this example, is controlled solely by the horizontal component of puck movements when the bottom button is held down.

4.4. Colour Selection Systems Evaluated

Twelve colour selection systems, each based on the colour order systems and interaction techniques described above, were evaluated in the colour matching experiment. A precise description of these selection systems is presented in Table 4.1. This table shows, for each selection system: the colour order system upon which the selection scheme is based, the interaction technique used, and how each of three degrees of freedom are controlled by the puck. For the RGB colour model, "Attribute 1" refers to the red coordinate, "Attribute 2" to the green coordinate, and "Attribute 3" to the blue coordinate. Similarly, attributes 1, 2 and 3 respectively refer to the hue, saturation and value properties of the HSV model. The Y, or luminance, channel in the YIQ system is attribute 1, while the I channel is attribute 2, and the Q channel is attribute 3. For the CIELAB system, the L (i.e. lightness), A and B channels are attributes 1, 2 and 3 respectively. The luminance property of the Opponent Colours system is referred to by attribute 1, the red-green property by attribute 2, and the yellow-blue property by attribute 3. The mnemonics given to each of the systems in Table 4.1 are used to reference the corresponding selection systems later in this thesis.

	Colour	Interaction	Puck Button Mappings		
Mnemonic	Model	Technique	Attribute 1	Attribute 2	Attribute 3
Rgb1	RGB	Single Attribute Control	Left Button, Horizontal Axis	Right Button, Horizontal Axis	Top Button, Horizontal Axis
Rgb2	RGB	Multiple Attribute Control	Top Button, Horizontal Axis	Top Button, Vertical Axis	Bottom Button, Horizontal Axis
Rgb3	RGB	Multiple Attribute Control	Top Button, Horizontal Axis	Bottom Button, Horizontal Axis	Top Button, Vertical Axis
Yiq1	YIQ	Single Attribute Control	Top Button, Horizontal Axis	Left Button, Horizontal Axis	Right Button, Horizontal Axis
Yiq2	YIQ	Multiple Attribute Control	Bottom Button, Horizontal Axis	Top Button, Horizontal Axis	Top Button Vertical Axis
Hsv1	HSV	Single Attribute Control	Left Button, Horizontal Axis	Right Button, Horizontal Axis	Top Button, Horizontal Axis
Hsv2	HSV	Multiple Attribute Control	Top Button, Horizontal Axis	Top Button, Vertical Axis	Bottom Button, Horizontal Axis
Lab1	CIELAB	Single Attribute Control	Top Button, Horizontal Axis	Left Button, Horizontal Axis	Right Button, Horizontal Axis
Lab2	CIELAB	Multiple Attribute Control	Bottom Button, Horizontal Axis	Top Button, Horizontal Axis	Top Button Vertical Axis
Opp1	Opponent Colours	Single Attribute Control	Top Button, Horizontal Axis	Right Button, Horizontal Axis	Left Button, Horizontal Axis
Opp2	Opponent Colours	Multiple Attribute Control	Bottom Button, Horizontal Axis	Top Button, Vertical Axis	Top Button Horizontal Axis
Opp3	Opponent Colours	Multiple Attribute Control	Top Button, Vertical Axis	Top Button, Horizontal Axis	Bottom Button, Horizontal Axis

Table 4.1. Description of the colour selection systems evaluated in the colour matching experiment.

Each of the five colour order systems was implemented once using the single attribute control technique and at least once using multiple attribute control. For the multiple attribute control technique, the "brightness" attributes, or "blueness" in the case of RGB, were singled out for independent control; that is, controlled by the puck's bottom button. The remaining two colour properties of each colour model were controlled by the vertical and horizontal puck movements when the top button was depressed. The decision to separate brightness was based on the intuition that this colour property is an independent attribute of the human colour perception process; isolating the colour blue in the case of the RGB model was a little more arbitrary. In addition to these systems, a selection system based on the Opponent Colours model, in which the yellow-blue axis was controlled by the bottom puck button and the top button simultaneously controlled the red-green and luminance axes, was also included in the evaluation. An implementation of the RGB model that combined the red and blue attributes for simultaneous control was also analyzed.

5. Experiment Description

In this chapter, a detailed description is given of the colour matching experiment upon which this initial investigation into interactive colour selection has been based. Topics include details of the hardware used, the physical environment in which the experiment was conducted, the task subjects performed, how subjects were selected, the instructions given to subjects, and what data was collected during the experiment.

5.1. System Details

The colour matching experiment was implemented using a VAX 11/780 host computer and an Adage-Ikonas RDS-3000 raster graphics system linked by a high speed DMA interface. The colour matching program was run on the host computer at a higher than normal priority so that the response of the system to the interactions of subjects performing the experiment would not be noticeably affected by the load caused by other users on the system. The graphics system displayed images using a 512 by 512 grid of picture elements and represented colour or depth information using 24 bit planes. To permit the colour of the test stimulus to be manipulated in real time, however, the stimulus was represented using only two pixel values (one for the fixed subfield and one for the variable subfield) and all colour manipulation operations were performed on the corresponding colour table entries. Colours were represented in the colour tables with 10 bits of precision for each of the red, green and blue primaries.

The matching stimulus was displayed on a 13-inch Electrohome ECM 1301 colour monitor. Both the brightness and contrast controls of the monitor were set midway between their minimum and maximum levels. The monitor was calibrated according to the procedures described in [Cowan83a] so that displayed colours could be accurately recorded in standard colour coordinates.

A Summagraphics Bit Pad One tablet equipped with a four button puck was used as the interaction device with which subjects controlled the colour of the test stimulus. The four coloured puck buttons were covered with white stickers so that subjects would not confuse the colours of the buttons with the colour properties the buttons controlled. In the instructions given to subjects, the buttons were referenced according to their position on the puck (i.e. left, right, top, bottom) rather than by their colour.

An Ann Arbour Genie video terminal was employed to display the instructions that guided subjects through the experiment. The display screen of this terminal used a neutral white phosphor to display alphanumeric characters. The contrast control of the terminal was set to its minimum level in order to minimize the brightness of the characters displayed on the screen.

5.2. Physical Environment

The experiment was conducted in a darkened room, the layout of which is shown in Figure 5.1. The setup illustrated was intended for right-handed subjects. Left-handed subjects used a symmetric layout.

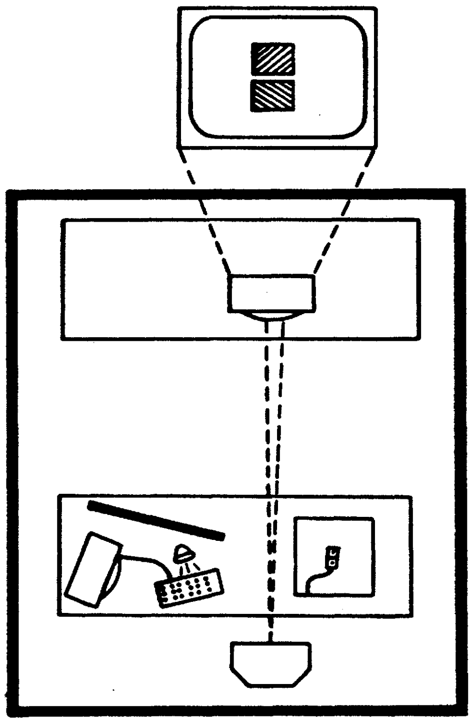

Figure 5.1. The layout of the room in which the experiment was conducted.

The colour monitor on which the test stimulus was displayed was positioned at about eye level, approximately 2.5 metres directly in front of the subject. This distance was necessary to prevent the subject from perceiving scanline effects. The test stimulus appeared in the form of two small rectangles. These rectangular subfields were positioned one above the other in the centre of the display monitor and were separated by a distance of 0.5 millimetres (two scanlines). Each subfield was 4.0 centimetres high and 5.0 centimetres wide and, hence, subtended approximately 0.9 degrees vertically and 1.2 degrees horizontally about the subject's direction of sight. Thus when the subject fixated on the matching stimulus, most of the light received by the subject's eye was focused on the fovea. It is in the fovea and the area immediately surrounding it that visual perception is most acute and the discrimination of colours is most sensitive [Wyszecki82].

The tablet and puck, the instruction terminal with its detached keyboard, a small desk lamp, and a partition were situated on a table directly in front of the subject. To minimize distraction during the colour matching session, the instruction terminal was placed in the periphery of the subject's field of vision on the side opposite the tablet. That is, for right-handed subjects, the tablet was placed to the right so that the right hand could be used for puck control and the terminal was placed to the left. The variable intensity desk lamp provided a very low level of incandescent illumination on the keyboard. The partition was strategically placed to avoid reflections of the lamp and terminal on the colour monitor. For the same reason, black felt was used to cover parts of the walls, ceiling and tables.

5.3. Experimental Procedure

Before actually beginning the experiment, each subject was requested by the experiment supervisor to read and sign a consent form and to complete a subject information form. A colour vision test (the Standard Pseudoisochromatic Plates test conducted under a fluorescent light source) was then performed to determine whether the subject suffered from any of the common forms of "colour blindness." If the test was passed, the subject was then shown to the room in which the experiment was to be performed and briefly introduced to the computer system. This short introduction consisted of informing the subject that the instructions for operating the tablet and for performing colour matching would be obtained by interacting with the alphanumeric terminal. (These instructions are discussed in Section 5.5). The experiment supervisor then left the room, leaving the subject to perform the experiment. The subject remained unaccompanied for the duration of the experiment. At the end of the experiment, the subject was asked to complete a comment form.

The task each subject was required to perform consisted of completing a total of 30 *trials*, or *colour matches*. All 30 matches were performed using the same colour selection system.

The procedure for each match was as follows. As is shown in Figure 5.1, the test stimulus used in this experiment consisted of two rectangular subfields vertically oriented on the display monitor. At the beginning of each match, the two subfields were set to different colours. The preselected colour of the top rectangle (the target colour) remained the same for the duration of the match. The colour of the bottom rectangle was controlled by the subject via the preselected colour selection system. The subject used the tablet and puck to interactively manipulate the controllable colour until it was perceptually equivalent to the target colour. Once a match had been achieved, the subject pressed a predesignated key on the keyboard to signal the completion of the match. The subject was then prompted to press the key again to begin the next match. Thus, the subject was allowed to rest between matches.

A maximum of three minutes was allowed to complete each match. If the subject timed out, a message indicating this was displayed on the instruction terminal and a prompt to start the next match was given. The allowed time interval appeared to be appropriate since it was long enough to allow the majority of matches to be completed without timing out and was short enough to urge subjects who were having difficulty (especially at the beginning of a session) on to subsequent matches.

Each subject matched the same set of target colours in the same order, regardless of the colour selection system used. Five colours were matched six times each. The standard CIE XYZ tristimulus values of the five target colours are specified in Table 5.1. Subjects were required to match this small set of colours several times rather than a set of 30 unique colours so that it would be possible to measure the learning component of task performance. The five target colours were chosen according to three criteria. First, each colour is representative of a different part of the chromaticity gamut. Second, each colour requires a nontrivial amount of effort on the part of the subject to achieve a match, regardless of the colour selection system used. Third, each colour is well within the displayable gamuts of all the colour order systems evaluated in the experiment. Colours situated near the boundaries of a colour system were avoided so that subjects would not frequently encounter these boundaries. The order in which the colours were matched is given in Table 5.2. This sequence was blocked into 6 different pseudorandom orders to eliminate the possibility that the performance of a match is affected by the colour of the previous match.

The initial colour of the controllable subfield (i.e. the colour from which the matching operation was begun) was the same for each trial, regardless of the target colour matched. The XYZ tristimulus values for the neutral grey starting colour are given in Table 5.1. The fact that the starting colour was always the same implied that for some target colours the match was begun further away from the target than for others. This was expected to have an insignificant effect on most performance measures since it was anticipated that most of the matching time would be spent making fine adjustments after the controllable colour was already in the proximity of the target colour. For criteria where the different distances between the starting colour and the target colour might prove to be significant, relative performance can still be measured in a meaningful way by comparing the data of a given target colour across selection systems.

	Descriptive	CIE XYZ Tristimulus Values		
Usage	Name	X	Y	Z
Starting Colour	Grey	0.08	0.08	0.07
Target Colours	Red	0.28	0.16	0.04
	Green	0.12	0.30	0.11
	Blue	0.09	0.05	0.33
	Yellow	0.24	0.28	0.06
	White	0.19	0.19	0.15

Table 5.1. Test stimulus colours used in the colour matching experiment, including the initial colour of the variable subfield and the set of five target colours that subjects were required to match.

Block 1:	1. Red	2. Green	3. Blue	4. Yellow	5. White
Block 2:	6. Blue	7. Red	8. Green	9. Yellow	10. White
Block 3:	11. Blue	12. Green	13. Yellow	14. White	15. Red
Block 4:	16. Blue	17. White	18. Green	19. Yellow	20. Red
Block 5:	21. Green	22. White	23. Red	24. Green	25. Yellow
Block 6:	26. Red	27. Blue	28. White	29. Blue	30. Yellow

Table 5.2. The order in which the target colours were matched. The order progresses left to right, top to bottom.

5.4. Subject Selection

The subjects that performed the experiment were partitioned into two principal categories according to the level of experience that individuals had in performing colour selection tasks and using interaction devices such as tablets. The first group consisted of subjects who had little or no experience with colour mixing techniques or interaction devices. Subjects that possessed a high degree of proficiency in performing colour selection or who had a lot of experience using interaction devices made up the second category. Artists, for example, who have spent years mixing paint colours would be considered to be part of the second group of "experienced" users.

One of the chief issues that were to be formally examined in the experiment was the issue of how subjects learned to perform the colour matching task using different colour selection systems. Thus, the number of inexperienced subjects that performed the experiment greatly exceeded the number of subjects from the experienced group. Each subject of the inexperienced category performed exactly one session and, therefore, used only one of the 12 colour selection systems. Subjects from this group were not given any initial description of the colour order system that they were to use for performing the colour matching tasks. Rather, they became familiar with the colour model by interacting directly with the computer graphics system (as is described in the next section). All the inexperienced subjects were volunteers, most of them from undergraduate computer science courses at the University of Waterloo.

The investigation into the performance of experienced subjects was carried out in a more informal manner. Each subject performed several sessions, each time using a different colour selection system. The analysis of task performance subsequently focused on how individual subjects performed using the different colour selection techniques rather than across subjects. While subjects from the inexperienced category became familiar with the colour selection system used through direct interaction, the experienced subjects had a conceptual understanding of the colour order system which they were to use before the beginning of each session. Thus, the learning factor of task performance was not expected to be as significant as with the inexperienced subjects. It was also anticipated that performance would not be significantly affected by the fact that a different colour selection system was used in each session. That is, the selection system used in the previous session was assumed to have a negligible impact on performance in the following session.

5.5. Subject Instruction

Prior to performing the experiment, subjects were familiarized with the computer graphics system and the colour matching operation through an interactive tutorial. This tutorial involved reading a total of 14 screens, or pages, of textual information and instructions, and performing 5 introductory practice "exercises." The pages of instructions were read in a sequential manner analogous to reading a book or manual. Subjects were able to return to previous pages, as well as paging forward, and could therefore reread directions as many times as they felt was necessary. Subjects were encouraged to seek the assistance of the experiment supervisor if the instructions were not thoroughly understood. The complete transcript of these instructions is given in Appendix B.

Subjects were trained for the experiment in an interactive fashion primarily because this method of instruction minimized the involvement of the experiment supervisor and consequently lowered the risk of biasing the experiment. It was also felt that this "hands on" method was the most appropriate way of demonstrating an interactive process such as colour matching.

Certain segments of the instructions, where the wording of the instruction's text is considered to have a significant effect on how subjects subsequently performed the colour matching operation, are now discussed. The instructions began with a brief introduction to the system followed by a description of the colour matching task. Specifically, to match colours, subjects were told to:

> ... manipulate the colour of the bottom square until it is the same as the colour of the top square.
>
> Note that the two colours should be matched as closely as possible. This means that you should continue in your attempt to match the colours until you think they are the same or until it becomes EXTREMELY difficult to get the colours any closer to each other. Each match should take you about 1 to 2 minutes to complete. You will actually be allowed 3 minutes for each match.

Subjects were informed of how long a typical match should take so that the amount of effort applied to achieve each match would be approximately uniform across subjects.

Following this description of colour matching, subjects were instructed how to manipulate the colour of the test stimulus. In order to emphasize the three dimensional nature of colour, an analogy to the length, width and height of a box was used. The following excerpt from the text of the instructions shows how this analogy was presented:

> Manipulating colour, as you will probably find out, is not as simple as you may have first thought. The best way to think of colour is to picture it as some three dimensional object such as a large box.
>
> A box has three properties, or dimensions: width, length and height. Like a box, colour also has three properties (or dimensions). Thus, there will be three properties of colour that you will be able to control in order to change the colour of the bottom square.
>
> If you like, you can imagine that the colour you are trying to match (ie. the colour of the top square) is a single point somewhere within the box. Your task then, is to search through the "contents" of the box and find the "right" point.

The remaining instructions, which included the 5 practice exercises, taught subjects how to use the puck and tablet in order to manipulate the three colour properties. The practice exercises were designed to incrementally introduce subjects to the interactive operations needed to perform the experiment. Emphasis was placed on effectively using the tablet and comprehending the colour matching operation rather than on becoming completely familiar with the behavior of the three colour properties. For instance, subjects were informed at one point:

> Don't get too worried if you don't understand the behavior of these properties right away - you will learn this as the experiment proceeds. After all, one of the goals of this experiment is to study this initial learning period!

The first practice exercise was merely intended to convince subjects that manipulating the puck in a particular way would indeed have an effect on the colour of the test stimulus. The second practice exercise tutored subjects on the use of the tablet in "mouse mode;" that is, using relative motion of the puck to cause colour alterations. These first two exercises also served the purpose of acquainting subjects with the first of the three colour properties that they controlled (e.g. brightness). The third and fourth exercises were intended as introductions to the other two controllable properties and instructed subjects on how these properties could be manipulated. The fifth and final exercise provided an opportunity for subjects to experiment with the three colour properties in combination and observe the effect on the colour of the test stimulus.

Immediately after completing the final practice exercise, subjects were informed that the actual experiment was about to begin. At this point, subjects could *not* return to reading the instructions. In fact, the subjects were not warned when they were at the end of the instructions in order to prevent them from practicing for an excessively long time. Furthermore, each practice exercise had a time limit of between 2 and 3 minutes. Subjects were generally discouraged from repeating these exercises, although they were permitted to do so if they thought it necessary.

The wording of the instructions did not rely on a subject's understanding of possibly unfamiliar colour terms. The property of "brightness," which is a common colour attribute that all subjects were thought to be familiar with, was referred to in the instructions in order to assist in demonstrating the interactions involved in performing the experiment. In the case of RGB, the concept of "redness" was used. For example, the subject was informed that if the puck was moved to the right while holding down a particular button, the colour of the comparison field would become brighter (or redder). In this way, the subject was able to determine if the instructions were being correctly followed since brightness (or redness) is a familiar concept. The characteristics of the two other colour properties with which the subject modified the colour of the test stimulus were not described in the text of the instructions; rather, the subject was briefly familiarized with the behavior of these properties through the interactive practice exercises. Thus, descriptive terms such as "hue," "saturation," and "tint" were avoided altogether.

5.6. Data Collection

Virtually all subject performance data was automatically generated and filed by the colour matching system. During the instruction phase of the experiment, the total duration of the introductory tutorial, the number of times each page of instructions was read, and the total durations of the 5 practice exercises were logged. For each match performed, the target colour, the final colour of the subfield controlled by the subject and the time taken to achieve the match were recorded. The latter measure is defined as the duration of the interval from the moment the test stimulus was initially displayed until the instant the subject signaled the completion of the match (or timed out). In addition to the above data, the colour of the controllable subfield was sampled once every second while the subject was performing the matching operation. This allowed the subject's path of colour manipulations to be traced at a later time. All changes in the status of the puck buttons were also logged. Such a status change occurred whenever a button was pressed and again when it was released.

In addition to the data generated by the system, information forms completed by subjects before and after the experiment were used to solicit comments, complaints, task experience levels and personal information. The subject information form provided a record of the subject's name, address, sex, age, department (if the subject was a student) and whether or not the subject's native language was English. The comment form completed at the end of the experiment is shown in figure 5.2. A summary of the comments that subjects recorded on this form is given in the next chapter.

COMMENT FORM

Have you ever had any prior experience with this kind of computer graphics equipment (i.e. the input tablet)? None ___ A Little ___ Lots ___

Have you ever had any experience mixing colours? Yes ___ No ___
If yes, please elaborate:

Did you have any difficulty understanding the instructions? Yes ___ No ___
If yes, please elaborate:

Any other comments on the experiment?

In order to help preserve the consistency of this study, we ask that you not discuss any *details* of the experiment with people who will be participating in the experiment at a future date (i.e. your classmates). Thanks again.

Figure 5.2. Information form on which subjects provided comments immediately after completing the experiment.

6. Results

Results of the colour matching experiment are presented in this chapter. A number of measures have been applied to the data collected during the experiment in order to quantify the behavior of the participating subjects according to several performance criteria.

The first three sections of this chapter deal with the performance of "inexperienced" subjects. In the first section, the performance measures are applied to the data separated by colour selection system. The primary goal in this section is to determine the relative performance of subjects using different colour selection techniques. In the second section, the issue of focus is "learning;" that is, how the performance of subjects changes with practice. Several performance measures are applied in such a way so as to determine the differences in learning between selection systems. The third section presents a summary of the written comments made by subjects at the end of each session.

The performance of "experienced" subjects is analyzed in the fourth section. Several of the performance measures used in the previous sections are also applied here to determine how well these subjects performed using the different selection systems.

All statistical measures and most of the graphs presented in this work were produced using the "S" data analysis software package [Becker84]. This analysis tool allowed all the relevant statistical tests to be calculated and graphically presented in a straightforward (and yet flexible) manner. The "S" package was chosen to assist in the analysis of the data produced in this investigation primarily because it appeared to be the most appropriate software tool that could run in a UNIX operating system environment.

Most of the data presented in this chapter is illustrated using bar graphs. In each bar graph, the height of the bars represent the *mean* value of performance. Each bar indicates the mean for a particular subset of the colour matching trials performed during the experiment. For instance, each bar in the graph may represent the mean value for all the colour matches performed using a particular colour selection system.

For many of the measures included in this analysis, 95 percent confidence intervals have been calculated for the means. These are indicated in the bar graphs by a vertical line extending through the centre of the top of each bar. The calculation of these confidence intervals and their meaning is discussed in [Ostle75]:

> Our interpretation of this statement [i.e. the confidence interval] is: the probability (before the sample is drawn) that the random interval (L, U), ... will cover or include the true population mean μ is equal to 0.95. Thus if a random sample is obtained from a normal population with mean μ and variance σ^2, and two quantities
>
> $$L = \overline{X} - t_{0.975(n-1)} s_{\overline{X}}$$
>
> and
>
> $$U = \overline{X} + t_{0.975(n-1)} s_{\overline{X}}$$
>
> are computed, we can say we are 95 percent confident that the true mean μ will be in the interval (L, U). [Note: $\overline{X}$ is the calculated sample mean, $s_{\overline{X}}$ is the standard error of the mean, and $t_{0.975(n-1)}$ is the 97.5 fractile of the Student's T distribution with n-1 degrees of freedom.]

Thus the confidence intervals should be interpreted as *a measure of certainty for the value of the mean*, not as an indication that 95 percent of the sample data had a value within the specified interval.

For those graphs where performance is shown by colour selection system, each bar is labeled using a mnemonic that indicates the selection system represented (e.g. Rgb1, Opp1, etc.). A description of the colour selection system that each mnemonic refers to is given in Table 4.1 in Chapter 4.

Many of the statistical measures applied in this analysis are based on the premise that the data sample has an approximately normal distribution [Ostle75]. While this assumption proved to be true for several performance measures, the samples for some measures were not distributed normally. For example, for measures of colour difference that indicate how close the colour of the two test stimulus subfields were at the completion of a trial, the samples tended to be positively skewed (i.e. away from zero). This is not surprising since the goal of the subjects was, after all, to *match* colours (i.e. achieve a colour difference of zero). In such cases, a logarithmic transformation was applied to the data in order to achieve a normal distribution. The means and corresponding confidence intervals were then calculated using the transformed data. In all cases, the transformation on the data was performed *before* the statistical measures were calculated. The Komogorov-Smirnov goodness of fit test was used throughout to measure the normality of the various subsets of data.

All measures of colour difference presented in this analysis were calculated using the CIE 1976 (L*a*b*) colour difference formula [CIE78]. This formula is one of two recommended by the CIE to predict "the magnitude of the perceived color difference between two given color stimuli." The value of colour difference between two colours calculated using this formula is equivalent to the Euclidean distance between the two colour values in the CIELAB colour space. Even though the actual colour matching task may have been performed using a colour order system other than CIELAB, it is felt that this difference measure is the most appropriate for quantifying colour differences since the CIELAB system is the closest existing approximation to a perceptually uniform colour system (see Section 4.2.4).

9.1. Performance of Inexperienced Subjects

The results presented in this section indicate the performance of inexperienced subjects. In total, 96 subjects from this experience category took part in the experiment. Thus there were 8 subjects (or 240 trials) for each of the 12 colour selection systems.

9.1.1. Task Incompletions

Subjects performing colour matching were allowed 3 minutes in which to complete the required task. Table 6.1 shows the number of times subjects exceeded this time limit by colour selection system.

The number of task incompletions can be interpreted as a measure of task difficulty. If subjects using one system "time out" more frequently than subjects using another, then this is an indication that subjects using the first system encountered greater difficulty in completing the task than the subjects using the latter system.

One pattern that can be seen in the information shown in Table 6.1 is that subjects who performed the colour matching tasks using selection systems based on the RGB or Opponent Colours models seemed to time out fewer times than subjects using other systems. In fact, these systems rank 6 of the "best" 7 in terms of this performance measure. On the other hand, subjects using the YIQ based systems exceeded the 3 minute time limit relatively frequently.

9.1.2. Trial Duration

The amount of time subjects spent performing colour matches is also a measure that is indicative of task difficulty. Figure 6.1 shows graphically how long subjects using the different colour selection systems took, on average, to complete the colour matching task. Trials where subjects did not complete the colour matching task within the 3 minute time limit were not included in calculation of this measure.

Subjects using the RGB or Opponent Colours based selection systems seemed to spend less time colour matching than subjects using systems based on other colour models. Again, these selection systems rate 6 of the best 7 according to this criteria. Subjects using the systems based on the HSV model in particular seemed to spend a relatively long time in performing the required tasks. As might be expected, these results are correlated somewhat with the number of task incompletions (as indicated in Table 6.1). The RGB and Opponent Colours systems tended to rank well as a group for both criteria, while the HSV, YIQ and CIELAB systems generally ranked lower.

9.1.3. Colour Difference

In Figure 6.2, the mean colour difference between the two colour matching stimuli at the point of task completion is graphed by colour selection system. A logarithmic transformation was applied to the measurements of colour difference in order to produce a normally distributed sample. Statistical calculations were then performed using this transformed data. Trials where

Task Incompletions	
System	Totals
Rgb1	11
Opp2	13
Opp1	14
Rgb2	14
Rgb3	15
Hsv1	19
Opp3	21
Lab1	23
Yiq1	34
Lab2	34
Hsv2	39
Yiq2	52
Total	289

Table 6.1. Number of times subjects failed to complete the colour matching task within the allowed time limit.

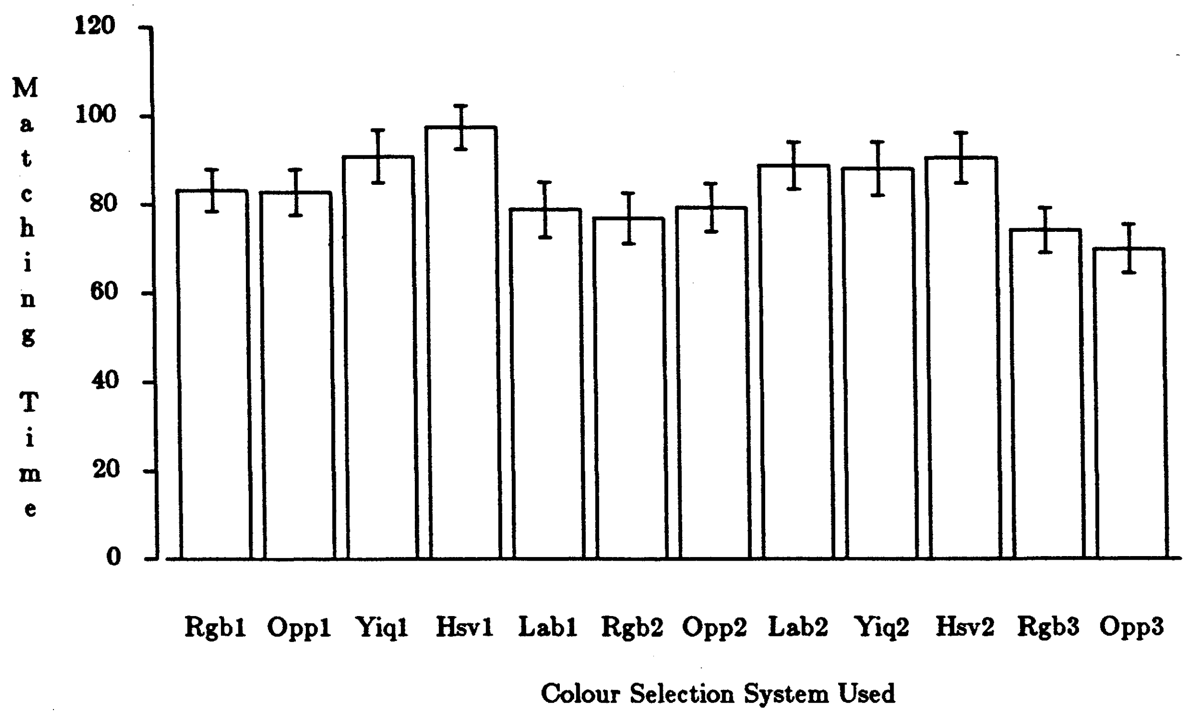

Figure 6.1. Mean time (in seconds) taken to complete the colour matching task. Incomplete trials are omitted.

subjects did not complete the colour matching task because the three minute time limit had been exceeded were not included in the calculation of these statistics.

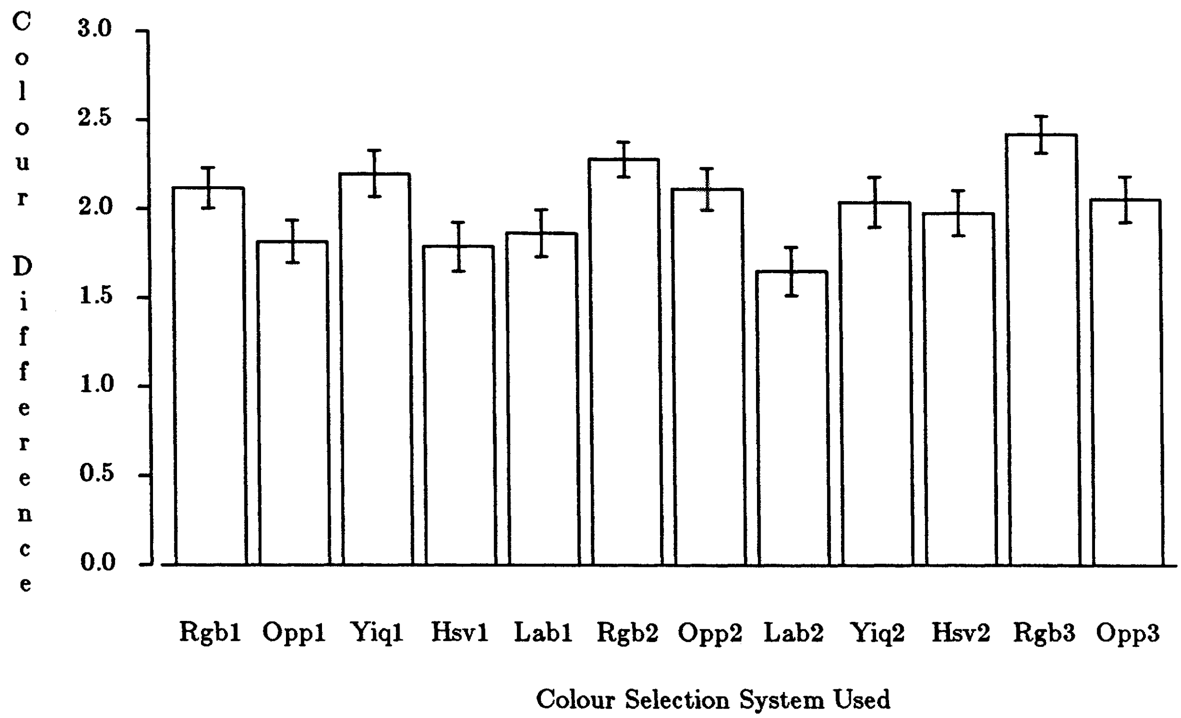

Figure 6.2. Mean log colour difference at completion of colour matching task. Incomplete trials are omitted.

The data illustrated in Figure 6.2 does not seem to contain any obvious patterns in terms of which colour order system or interaction technique appears to be best suited for the criteria of minimizing colour difference at task completion. However, it does seem that subjects using the selection systems based on the CIELAB colour model were able to match slightly closer than subjects using most of the other systems. In particular, the CIELAB based colour selection systems rated first and fourth for this performance measure; however, the differences in the mean values between the best 4 selection systems are not significant. At the other end of the spectrum, subjects using the RGB based systems appeared to perform considerably worse than subjects using other systems. The RGB systems ranked 3 of the worst 4 according to this criteria. A discussion on possible explanations for why subjects using RGB based systems seemed to complete the colour matching task quickly but did not match as closely as subjects using other systems is given in Chapter 7.

6.1.4. Hue, Chroma and Lightness Differences

In this section, the colour difference at the time of task completion (i.e. the measure presented in the previous section) is decomposed into three characteristic components: "hue," "chroma" and "lightness." Each of these perceptual properties are calculated according to the standard CIELAB based formulas as recommended by the CIE in [CIE78]. In this publication, it is recommended that these formulae be used "when it is desired to identify the components of color difference in terms of approximate correlates of lightness, perceived chroma, and hue." These difference measures are related to the integral CIELAB colour difference measure by the formula:

$$\Delta E = \sqrt{(\Delta H)^2 + (\Delta C)^2 + (\Delta L)^2}$$

where ΔE is the measure of integral colour difference and ΔH, ΔC, and ΔL are the difference measures of hue, chroma and lightness respectively. Thus, each component difference measure affects the value of the integral colour difference measure equally. Although these measures of colour difference are based on the CIELAB colour order system, the reader should be aware that the above attributes of hue and chroma are substantially different from the "A" and "B" axes of the CIELAB colour system (i.e. the colour properties subjects controlled during the experiment).

In Figure 6.3a, the mean difference in hue at trial completion is graphed by colour selection system used. Similarily, differences in chroma and luminance are plotted in Figures 6.3b and 6.3c. A logarithmic transformation was applied to each of these measures before calculating the means in order to achieve normally distributed samples. In addition, trials where subjects failed to complete the colour matching task in the alloted time were omitted in these calculations.

For the chromatic attributes of hue and chroma, there did not appear to be a clear indication of which colour order systems or interaction techniques were best suited for matching closest in terms of these properties. However, it is evident that the RGB and YIQ based systems did not rate very well for these measures. In particular, all of the RGB systems were among the worst 5 according to the criteria of minimizing chroma difference.

For lightness differences there are some significant differences in performance between colour selection systems. Performance generally varied more between selection systems in terms of lightness difference than for hue and chroma differences. In particular, subjects using the RGB based systems matched considerably less precisely in terms of lightness than subjects using other systems. Of interest here is the fact that the RGB systems were the only systems used in the experiment that did not have a separate degree of freedom for directly controlling lightness or a colour property that was perceptually quite similar. Also, it appears that this substantial disadvantage in performance in terms of lightness difference is largely responsible for the poor rating of the RGB systems for the integral colour difference criteria discussed in the previous subsection, although the RGB systems also rated poorly according to the hue and chroma difference measures.

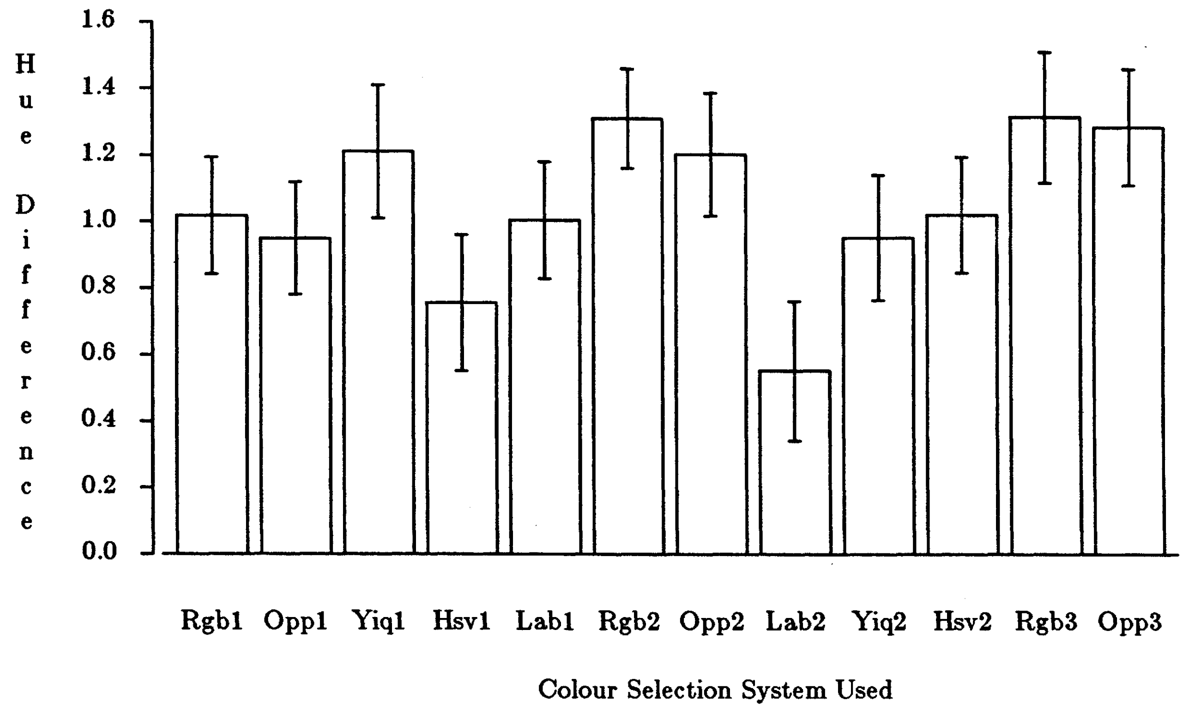

Figure 6.3a. Mean log hue difference at completion of colour matching task. Incomplete trials are omitted.

6.1.5. Colour Difference Thresholds

The performance of subjects is analyzed here in terms of several *thresholds* of colour difference. Each such threshold is a level of colour matching precision defined in CIELAB colour difference units. In colour matching, the difference between the two colours to be matched is initially relatively large. As the subject performs the colour matching task, the difference between the two colours is diminished (with varying degrees of fluctuation) until a subjective match is achieved. The point in a colour matching trial where the difference in colour between the two subfields is less than or equal to a particular threshold value for the *first* time is defined here as the point in the trial where the threshold is "reached."

The primary objective in analyzing subject performance using colour difference thresholds was to allow performance to be examined at different stages of task completion. In particular, it was of interest to determine whether subjects tended to perform in a relatively uniform manner while matching colours or whether performance seemed to change as subjects came closer to achieving a match.

Two primary measures were calculated for each threshold value: the proportion of trials for which the given threshold was reached and the mean time taken by subjects to reach the threshold. The times taken to reach the various thresholds were originally recorded in seconds. However, a logarithmic transformation was applied to this data since it was determined that this was necessary to obtain a sufficiently normal sample. The mean times taken to reach a threshold

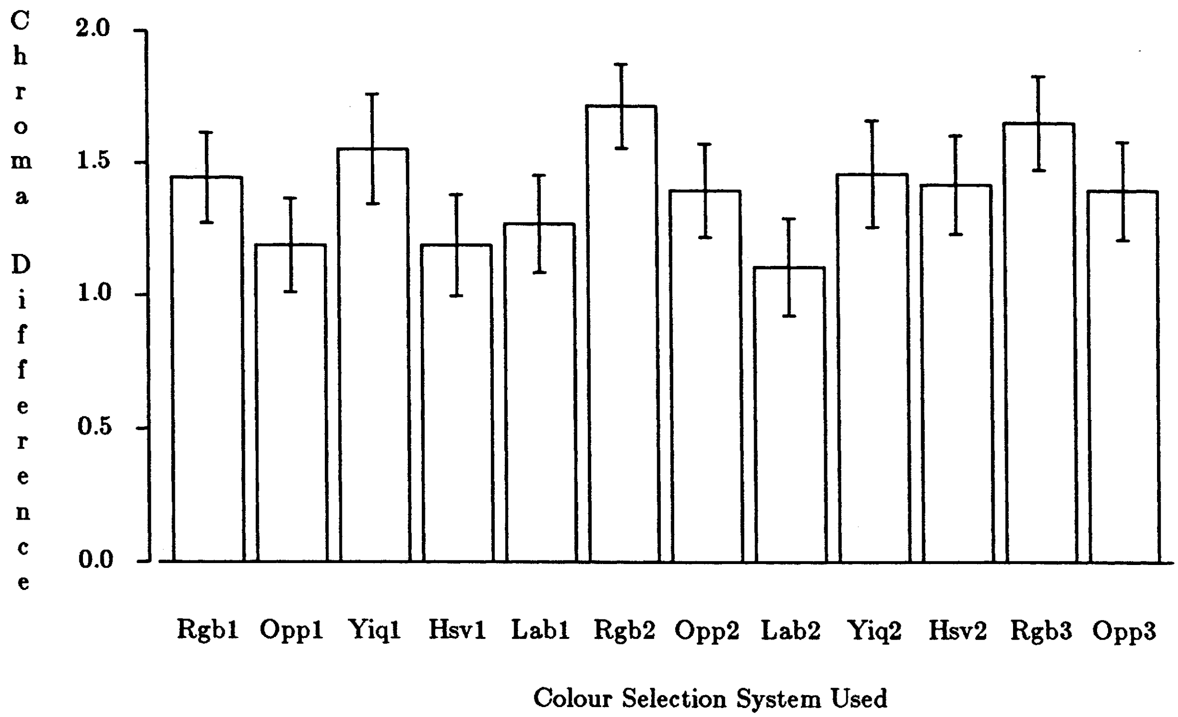

Figure 6.3b. Mean log chroma difference at completion of colour matching task. Incomplete trials are omitted.

were then calculated using the transformed data. The reader should keep in mind that these measures were based on transformed data when comparing to other measures of time presented in this thesis that were not calculated using transformed data (e.g. mean trial duration).

Initially, the measures described above were calculated for a total of 15 threshold values ranging between 2 and 40 CIELAB colour difference units. This data was examined to determine a small number of thresholds that would serve to represent the characteristic behavior of the subjects. A more rigorous analysis was then performed using these thresholds. The three threshold values that were chosen for this purpose are 25, 14 and 6 CIELAB colour difference units. The analysis performed using these values is presented in the remainder of this section.

Since the "white" target colour used in the experiment was approximately 15 CIELAB colour difference units from the starting colour of the variable subfield, all trials where the white colour was matched were omitted in the calculation of performance measures for threshold values greater than 15 CIELAB units (i.e. for the 25 CIELAB unit threshold). However, in order that meaningful comparisons can be made between performance measures calculated for threshold values greater than 15 CIELAB units and those calculated for threshold values less than 15 CIELAB units, each measure is calculated twice for the threshold values less than 15 CIELAB units: once including trials where white was matched and once with the white matches omitted. Trials where the colour matching task was not completed because the subject exceeded the allowed time limit *are* included in the calculation of these performance measures.

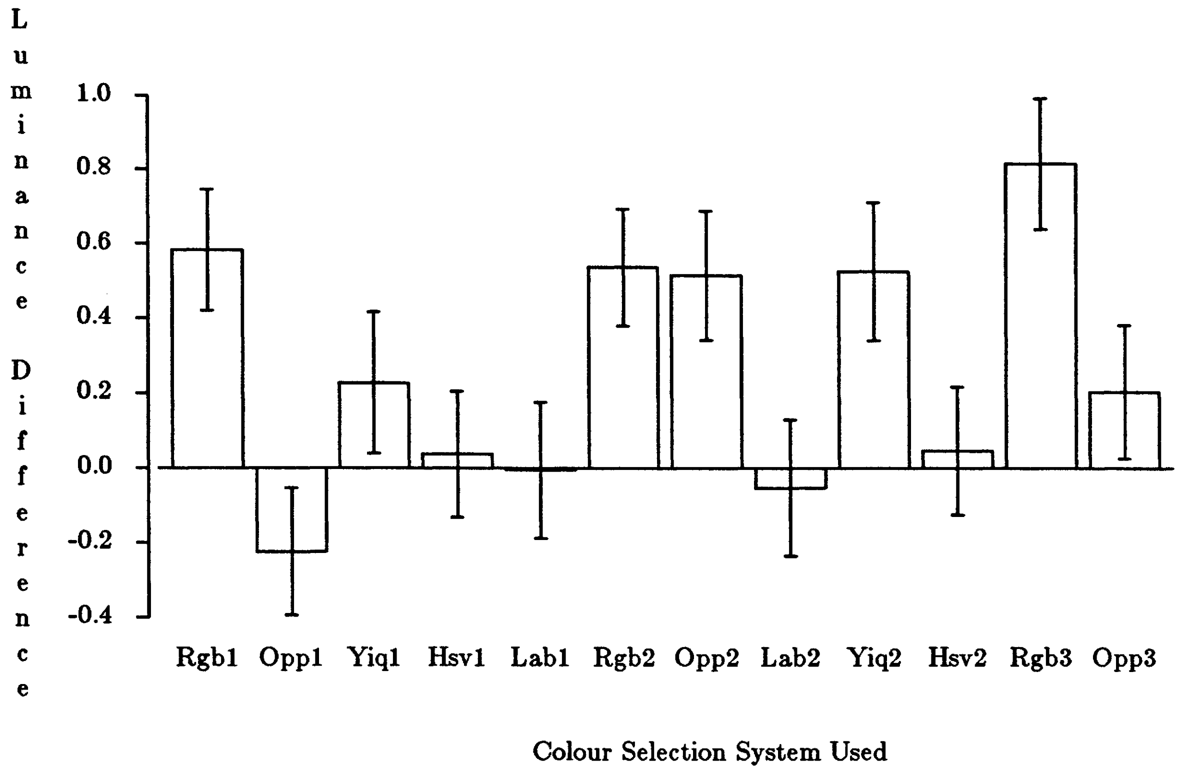

Figure 6.3c. Mean log luminance difference at completion of colour matching task. Incomplete trials are omitted.

The performance measures are presented here in figures that each contain two graphs. The upper graph in each figure displays the proportion of trials in which the threshold was reached. The lower graph shows the mean values of the time subjects took to reach the threshold.

The performance measures for a threshold value of 25 CIELAB units are graphed in Figure 6.4. In Figures 6.5a and 6.5b, the measures for a threshold value of 14 CIELAB units are displayed. Trials where white was matched were included in the calculations of the measures in Figure 6.5a, but were not included in the calculations for the measures shown in figure 6.5b. Similarly, figures 6.6a and 6.6b show the performance measures (including and excluding the white matches respectively) for a threshold value of 6 CIELAB difference units.

At a threshold value of 25 CIELAB units, it can be seen from the data graphed in Figure 6.4 that subjects who used selection systems based on the RGB or Opponent Colours models tended to reach the threshold faster than subjects using selection systems based on other models. In fact, the RGB and Opponent Colours systems ranked as the best 6 according to this criteria. Subjects using the HSV systems seemed to take considerably longer to reach the 25 CIELAB unit level than subjects using any of the other selection systems. There were no really significant differences between systems for the proportion measure. For almost every selection system, subjects were able to reach the 25 CIELAB unit threshold in over 90% of the trials.

At the 14 CIELAB unit threshold, the situation, according to the data displayed in Figure 6.5b, was much the same. The RGB and Opponent Colours based systems again ranked as the best 6 and the HSV systems rated 2 of the worst 3 in terms of the mean time taken to reach the threshold. The magnitude of the differences between systems here is, however, considerably less than at the 25 CIELAB unit level. Again, there were no major differences between systems according to the proportion measure. Subjects were able to reach the 14 CIELAB unit threshold in over 75% of the trials using any of the selection systems.

When the white matches are included at this threshold level, the RGB based systems rank significantly lower (according to Figure 6.5a), while the Opponent Colours systems still rate as the best (with the exception of "Opp2"). In addition, the HSV systems are still the worst in terms of this measure. The differences between systems in Figure 6.5a are, however, much lower than in Figure 6.5b. The proportion measure in Figure 6.5a is similar to that in Figure 6.5b but again the variances between systems are less significant.

At a threshold value of 6 CIELAB units, the differences between systems in the time taken by subjects to reach the threshold are, according to Figure 6.6b, very small compared to the variances at previous levels. Hence, there is no clear indication of which systems rank "best" for this criteria. The HSV systems again rank as 2 of the worst 3, but the significance of these rankings are minimal. In addition, there are not really any consistent patterns in the proportion measure at this level except perhaps that the RGB based systems rate rather poorly.

When the white matches are included at the 6 CIELAB level (see Figure 6.6a), the only notable difference from the results of Figure 6.6b is that the RGB based systems rate considerably worse, ranking 3 of the worst 4 for both performance measures.

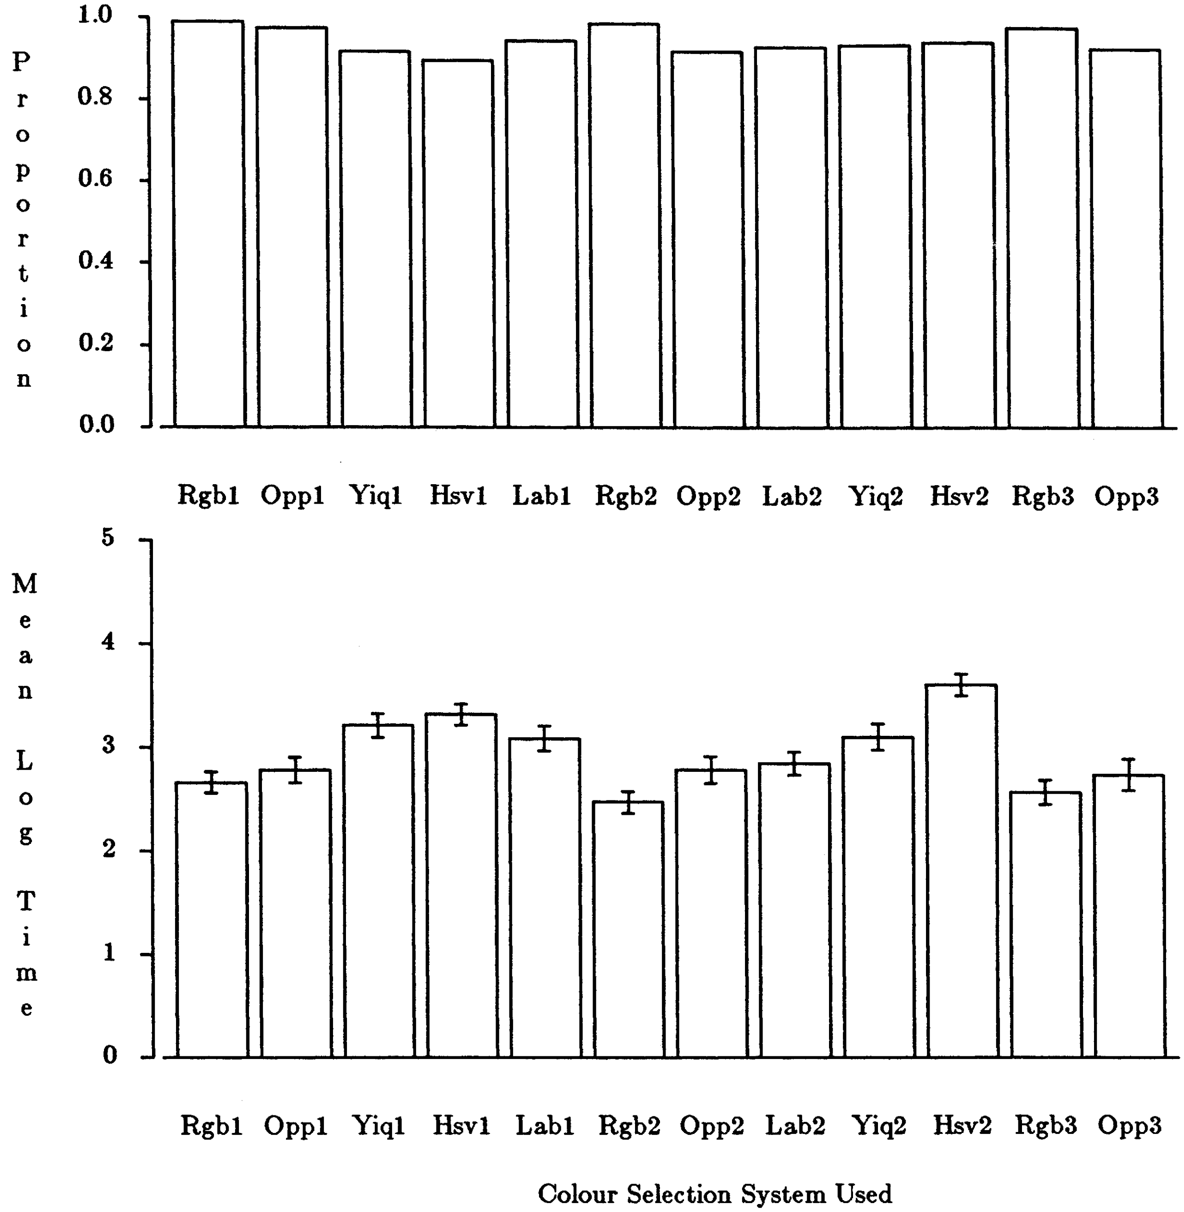

Figure 6.4. Threshold of 25 LAB colour difference units. Upper graph shows proportion of trials for which the threshold was reached. Lower graph shows mean log time taken to reach the threshold. Trials where white was matched are excluded from these measures.

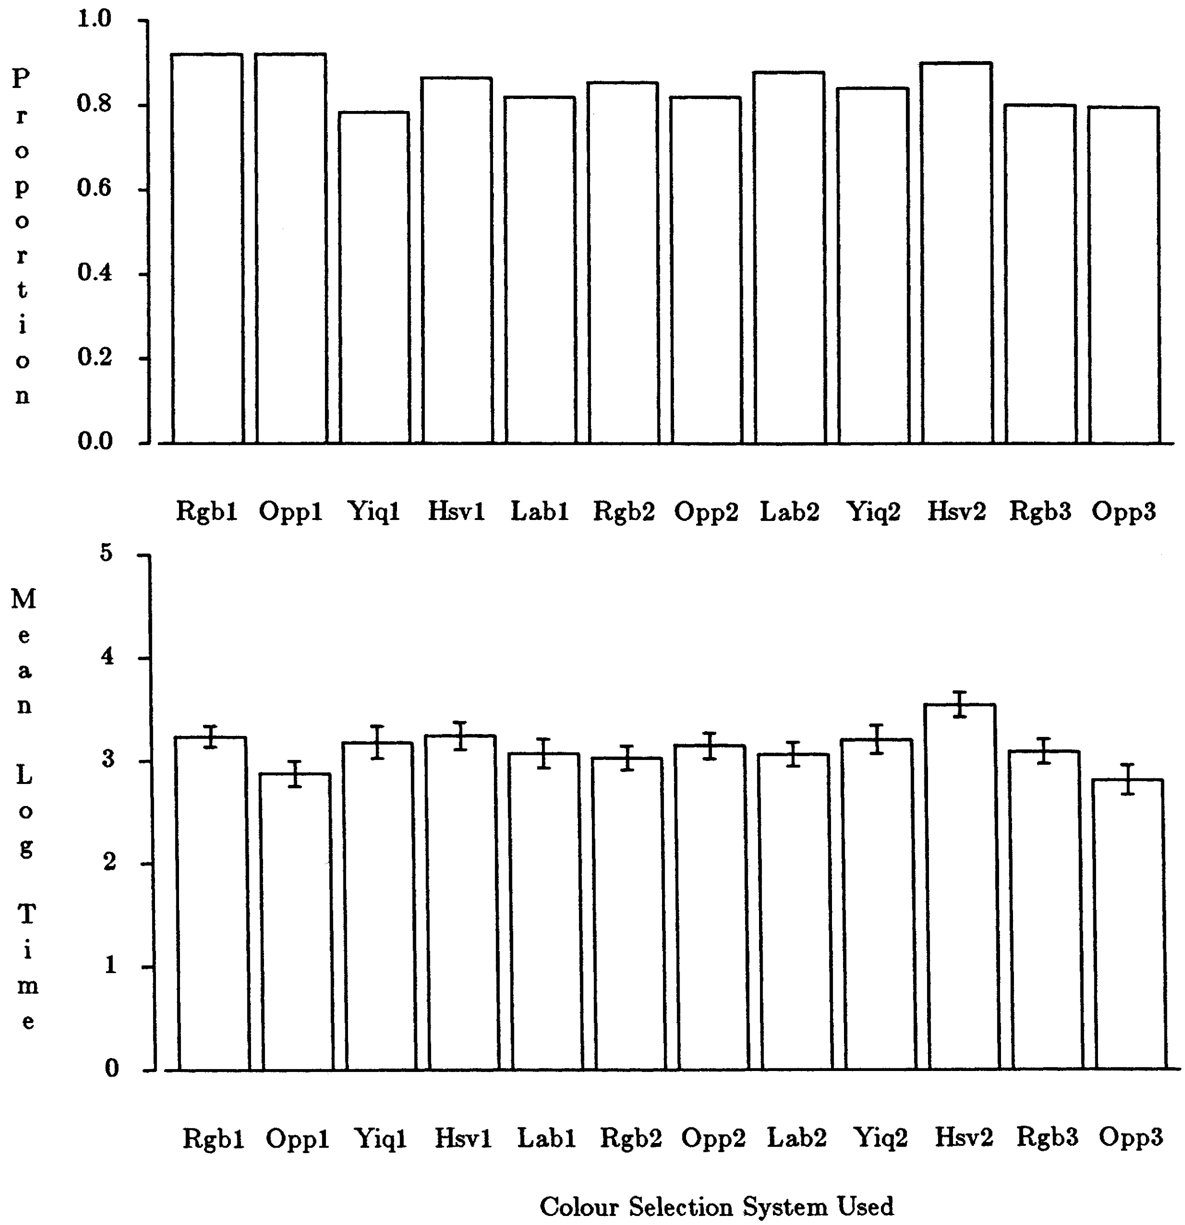

Figure 6.5a. Threshold of 14 LAB colour difference units. Upper graph shows proportion of trials for which the threshold was reached. Lower graph shows mean log time taken to reach the threshold.

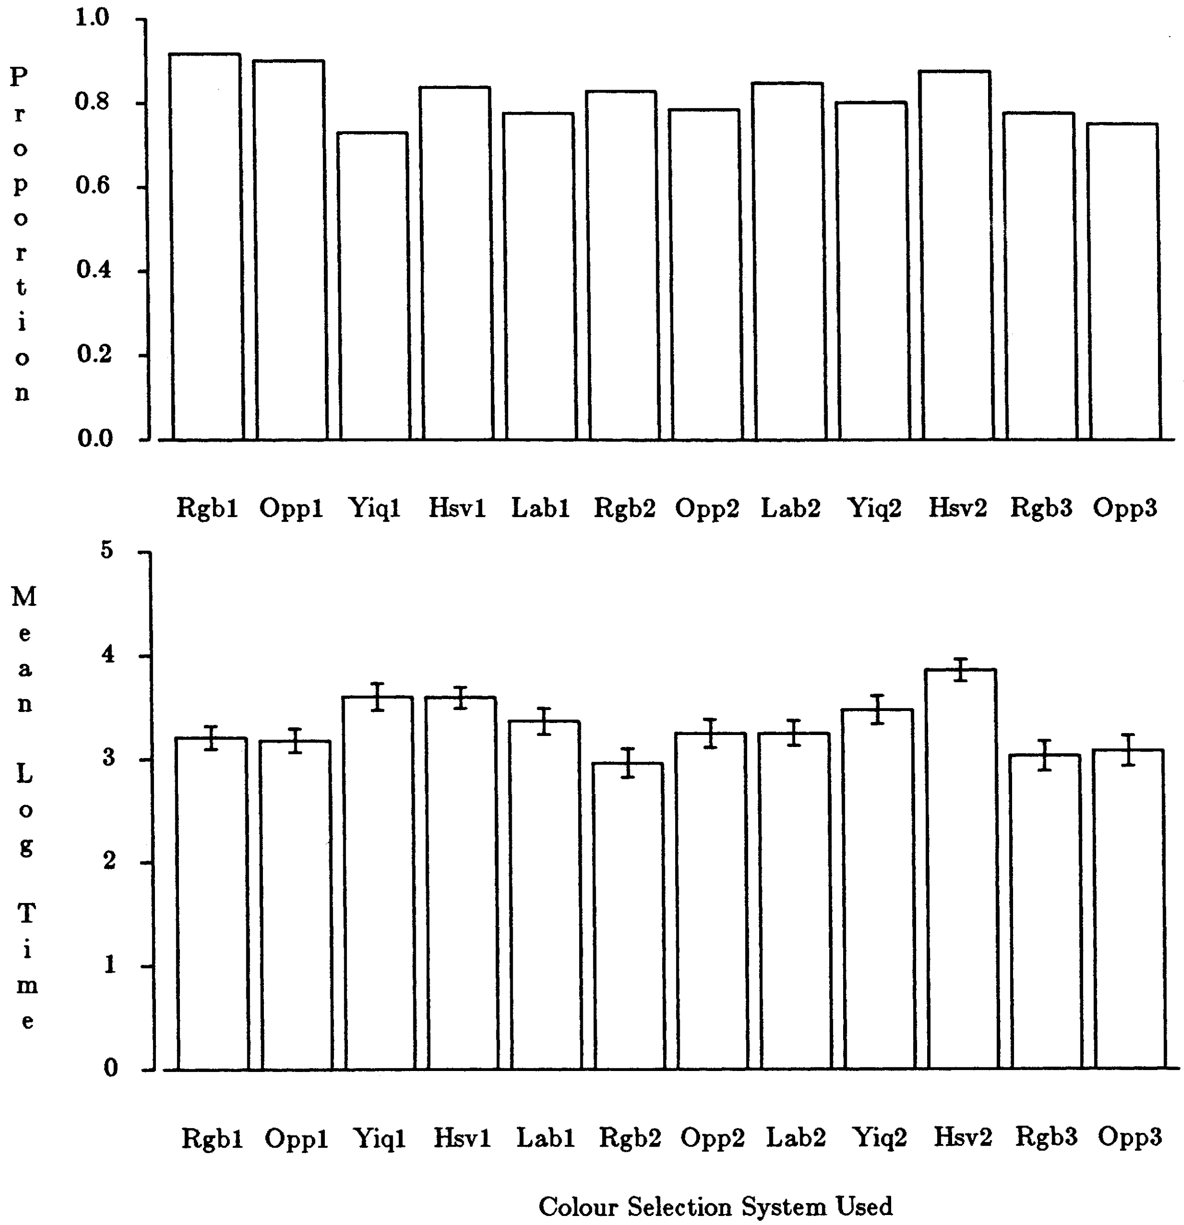

Figure 6.5b. Threshold of 14 LAB colour difference units. Upper graph shows proportion of trials for which the threshold was reached. Lower graph shows mean log time taken to reach the threshold. Trials where white was matched are excluded from these measures.

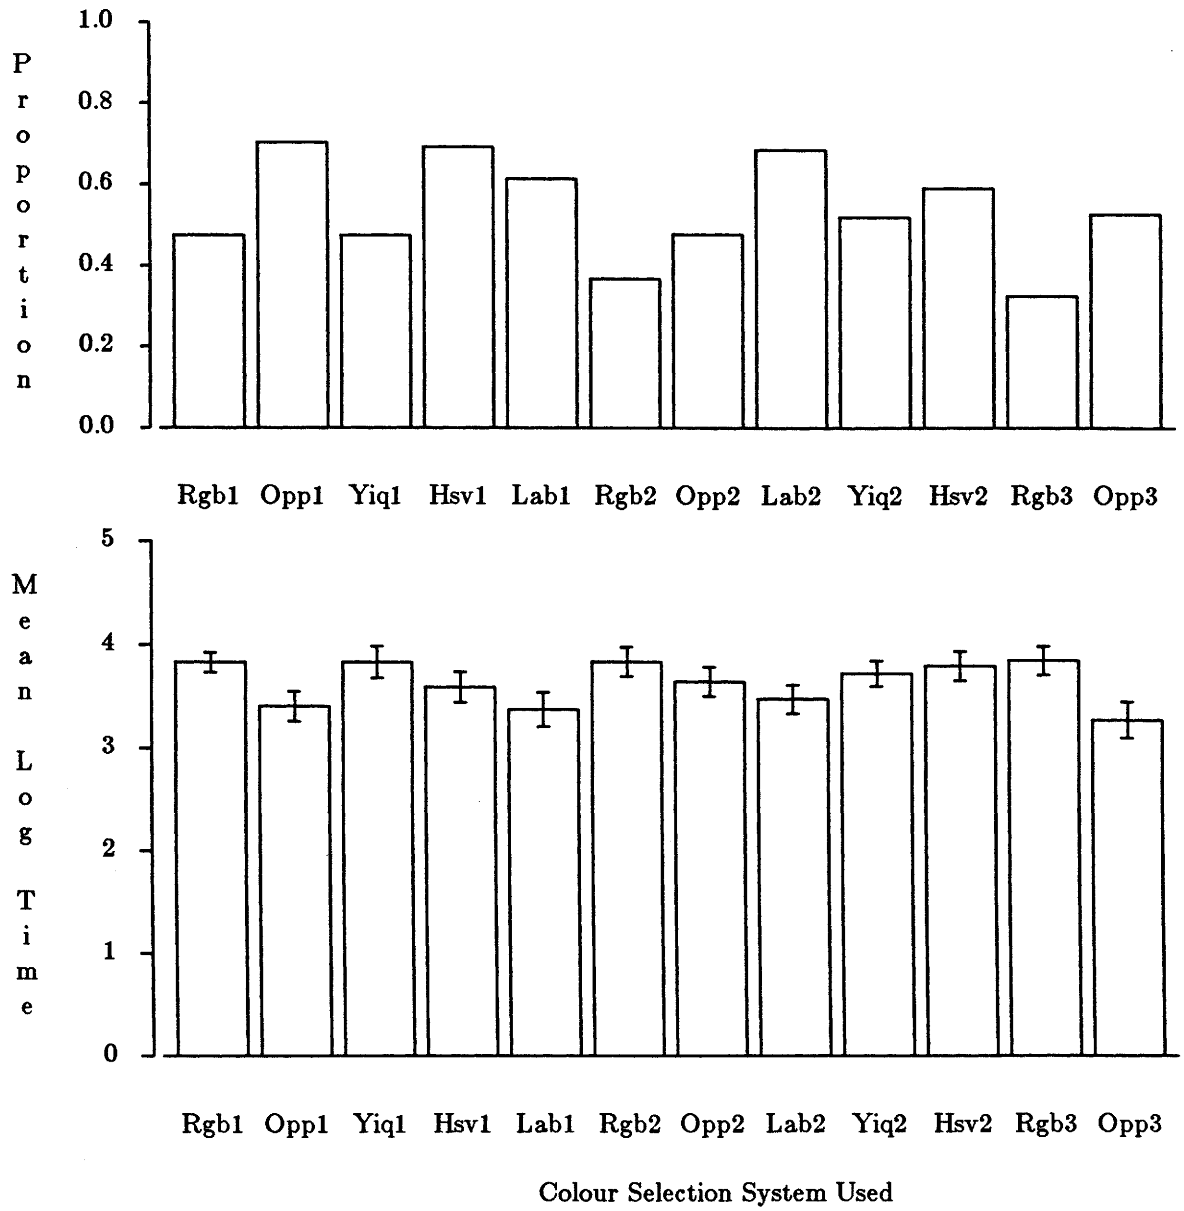

Figure 6.6a. Threshold of 6 LAB colour difference units. Upper graph shows proportion of trials for which the threshold was reached. Lower graph shows mean log time taken to reach the threshold.

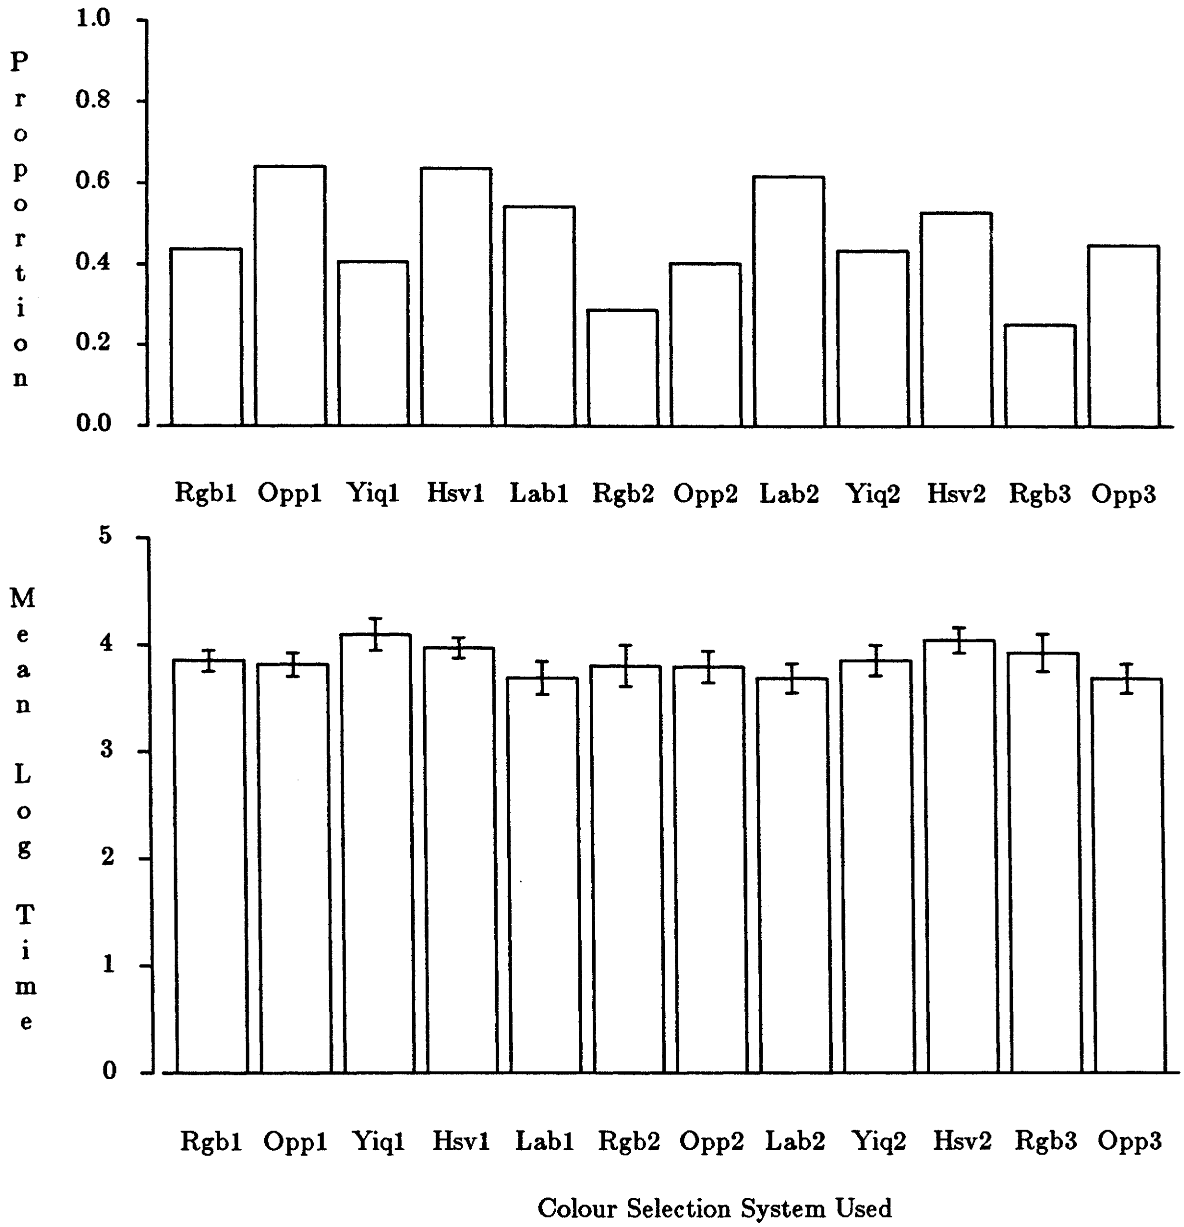

Figure 6.6b. Threshold of 6 LAB colour difference units. Upper graph shows proportion of trials for which the threshold was reached. Lower graph shows mean log time taken to reach the threshold. Trials where white was matched are excluded from these measures.

6.2. Learning

In this section, the learning component of subject performance is analyzed. A number of measures are used here to gauge the changes in performance as subjects became more familiar with the colour matching process and the colour selection system they were using. Each measure is calculated across all trials by "colour slice" and for each colour selection system by "session half." The Nth *colour slice* is defined as the set of trials within a session where a colour is matched for the Nth time. For example, the first colour slice consists of the first through fifth trials of each session - the first time a subject matches each of the 5 colours. The first *session half* simply consists of the first 15 trials in a session, while the second session half consists of the last 15 trials (of a total of 30).

When performance is compared across colour selection systems, the performance measures shown were calculated by session half rather than by colour slice. This was done in order to increase the statistical significance of the measures. That is, the sample size for each selection system would have been 40 had the data been split by colour slice. By grouping the data by session half instead, the size of each sample was increased to 120. It was determined that this substantially improved the significance of the calculated measures.

6.2.1. Task Incompletions

Table 6.2 shows, by colour slice and colour selection system, the number of trials where subjects did not complete the colour matching task within the allowed 3 minute time limit.

Task Incompletions							
	Colour Slice						
System	1	2	3	4	5	6	Totals
Rgb1	4	4	0	2	0	1	11
Opp2	5	3	2	2	0	1	13
Opp1	5	3	2	0	2	2	14
Rgb2	4	4	3	0	1	2	14
Rgb3	5	3	0	1	4	2	15
Hsv1	10	5	2	1	1	0	19
Opp3	6	4	3	2	1	5	21
Lab1	6	4	6	3	1	3	23
Yiq1	6	11	7	3	4	3	34
Lab2	10	4	7	8	2	3	34
Hsv2	8	12	6	5	4	4	39
Yiq2	14	11	8	7	8	4	52
Totals	83	68	46	34	28	30	289

Table 6.2. Number of times subjects exceeded the allowed time limit by colour slice and colour selection system.

From the totals shown in the bottom row of this table, it can be seen that there was a noticable decrease in task incompletions as subjects attained more practice. Most of this "learning" seemed to be concentrated near the beginning of each session since the largest decreases in total timeouts occurred between the first three colour slices. The data displayed in the table does not seem to show any obvious correlation between the rates of decrease in incompletions and the colour selection system used. For the most part, the data for each selection system shows a decrease in incompletions across colour slices, but the rates of decrease vary considerably from system to system and from colour slice to colour slice.

6.2.2. Trial Duration

Figure 6.7a shows the average time taken to complete the colour matching task by colour slice across all data, while Figure 6.7b shows this performance measure by colour selection system. In Figure 6.7b, the left (unshaded) bar of each pair of bars indicates performance during the first half of each session. The right (shaded) bar of each pair represents subject performance during the second half of each session. Trials where subjects did not complete the colour matching task because the three minute time limit was exceeded were not included in calculations of these measures.

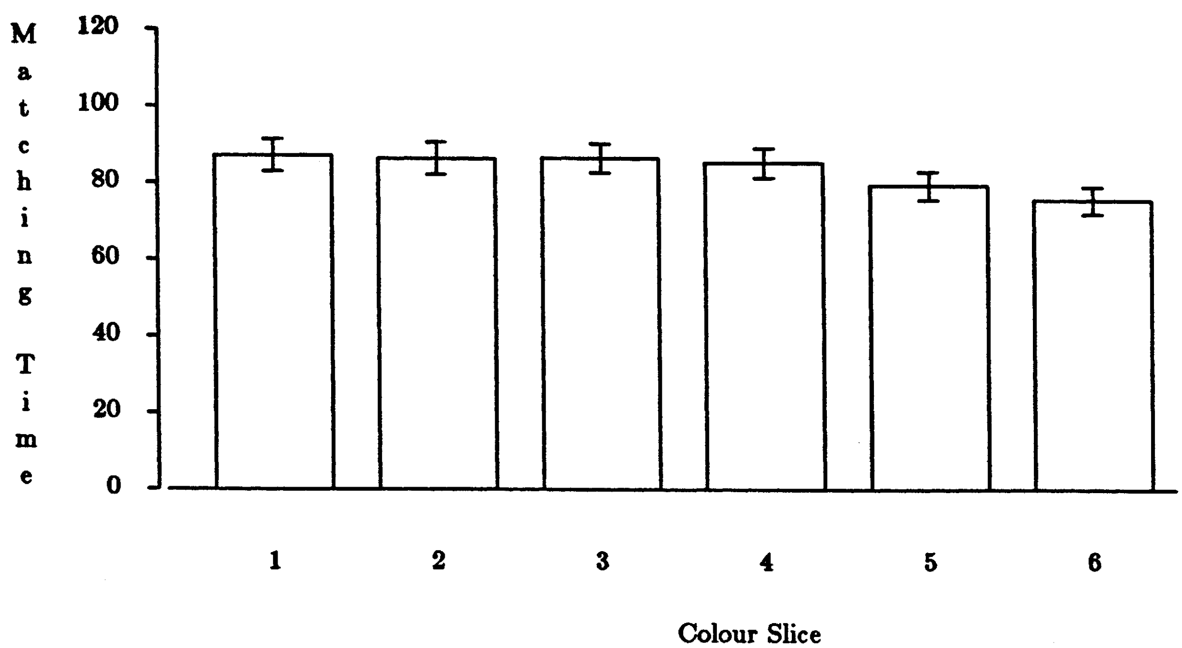

Figure 6.7a. Mean time taken to complete colour matching task by colour slice across all trials. Incomplete trials are omitted.

Figure 6.7a shows that the mean time taken by subjects to complete the colour matching task remained almost constant across the first four colour slices but then dropped considerably for the fifth and sixth colour slices. In Figure 6.7b, it can be seen that for most selection systems there is a noticable decrease in mean trial duration between the first and second session halves for most of the selection systems. There does not, however, appear to be any consistent pattern here in terms of which colour order systems and which interaction techniques exhibited considerably improved performance between session halves and which did not.

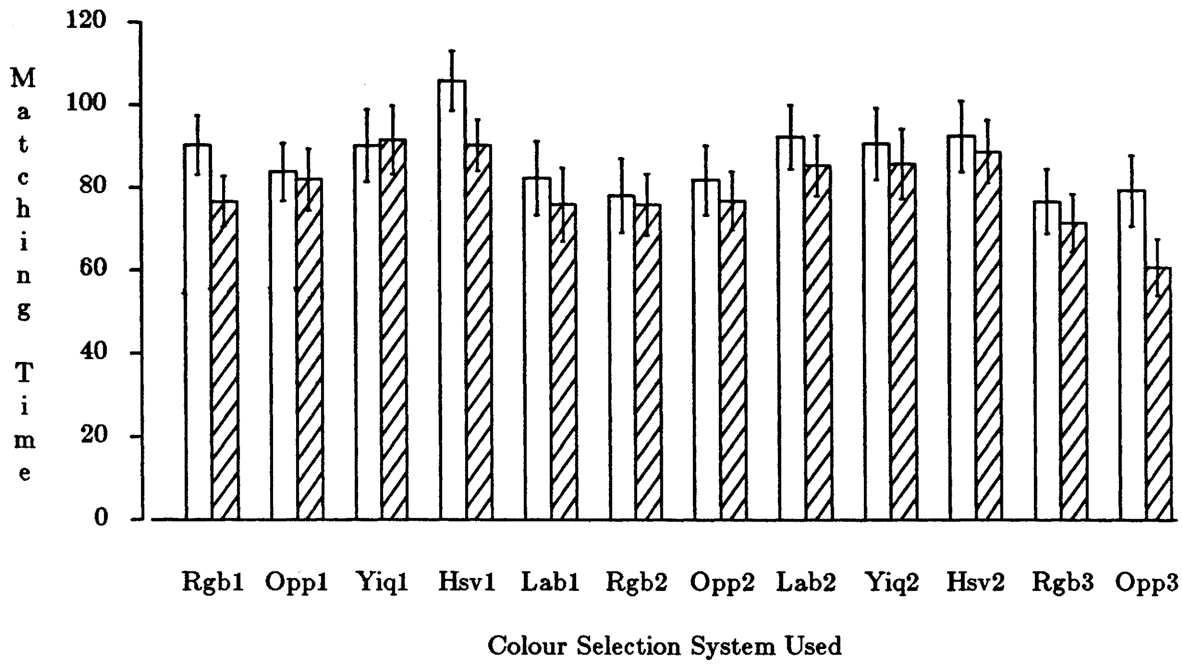

Figure 6.7b. Mean time taken to complete colour matching task by colour selection system and session half. The unshaded bars represent the means of those trials occuring in the first half of each session, while the shaded bars represent performance during the second half of each session. Incomplete trials are omitted.

6.2.3. Colour Difference

In Figures 6.8a and 6.8b, the average colour differences (between the colour of the variable subfield and the target colour) at the time of task completion are displayed. Figure 6.8a shows the colour difference measure calculated by colour slice across all trials. Figure 6.8b shows the mean colour differences by colour selection system and session half. A logarithmic transformation was performed on all the colour difference data presented here in order to map the data into a normally distributed sample. Again, trials where subjects did not complete the colour matching task in the allowed time were omitted from the calculations of these performance measures.

Figure 6.8a shows that there was a relatively large decrease in the mean colour difference measure between the first and second colour slices, followed by a rather small but steady decline across the next four colour slices and a sudden peculiar jump at the last colour slice. It is apparent then that most of the learning in terms of this measure occured right at the beginning of the session. Performance then improved only slightly for the remainder of the session. The increase in colour difference between the fifth and sixth colour slices may have occurred as a result of subjects becoming tired or anxious to finish as they neared the end of the experiment.

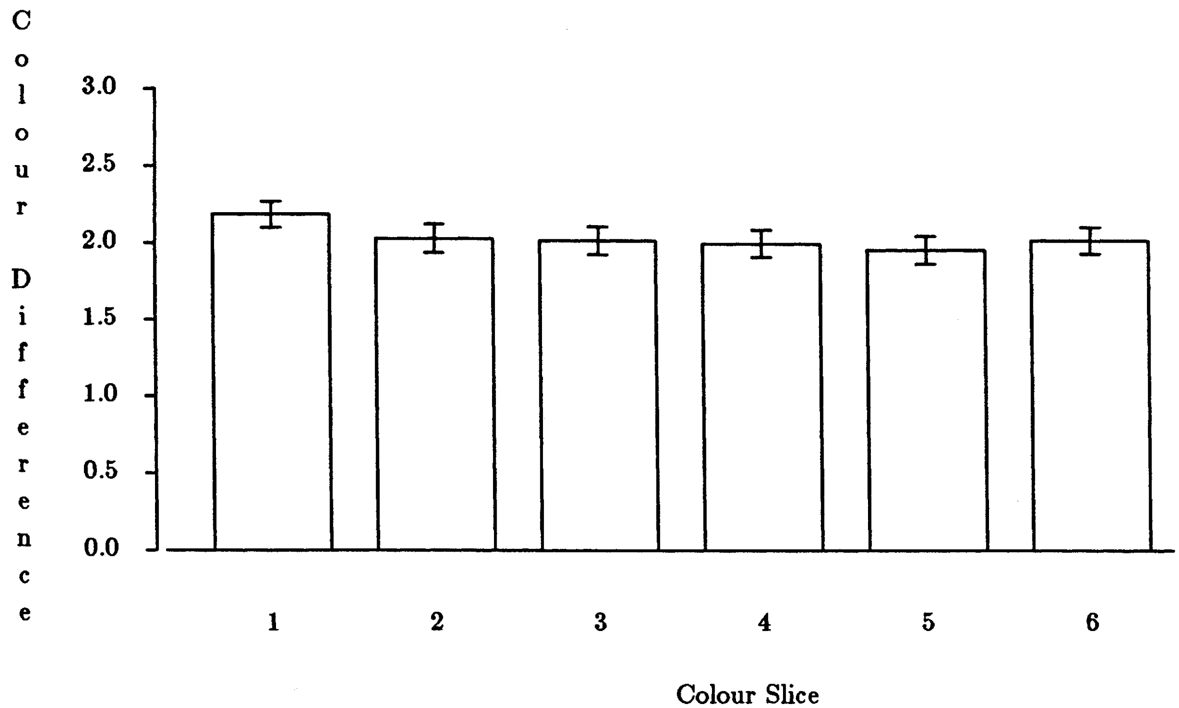

Figure 6.8a. Mean log colour difference at task completion by colour slice across all trials. Incomplete trials are omitted.

In Figure 6.8b, there is one pattern that is evident from the data shown. For each of the selection systems based on RGB or Opponent Colours, with the exception of the "Rgb2" system, there was little or negative improvement in the colour difference measure between the first and second session halves. In fact, these systems ranked as 5 of the lowest 6 in terms of order of improvement between session halves. All of the other systems evaluated showed a moderate decrease in mean final colour difference between session halves. The magnitude of these improvements in relation to the size of the confidence intervals, however, is rather small.

6.2.4. Colour Difference Thresholds

Changes in performance at several thresholds of colour difference are examined here. Again, a logarithmic transformation has been applied to the data to achieve a normal distribution. Incomplete trials *are* included in the calculations of these measures.

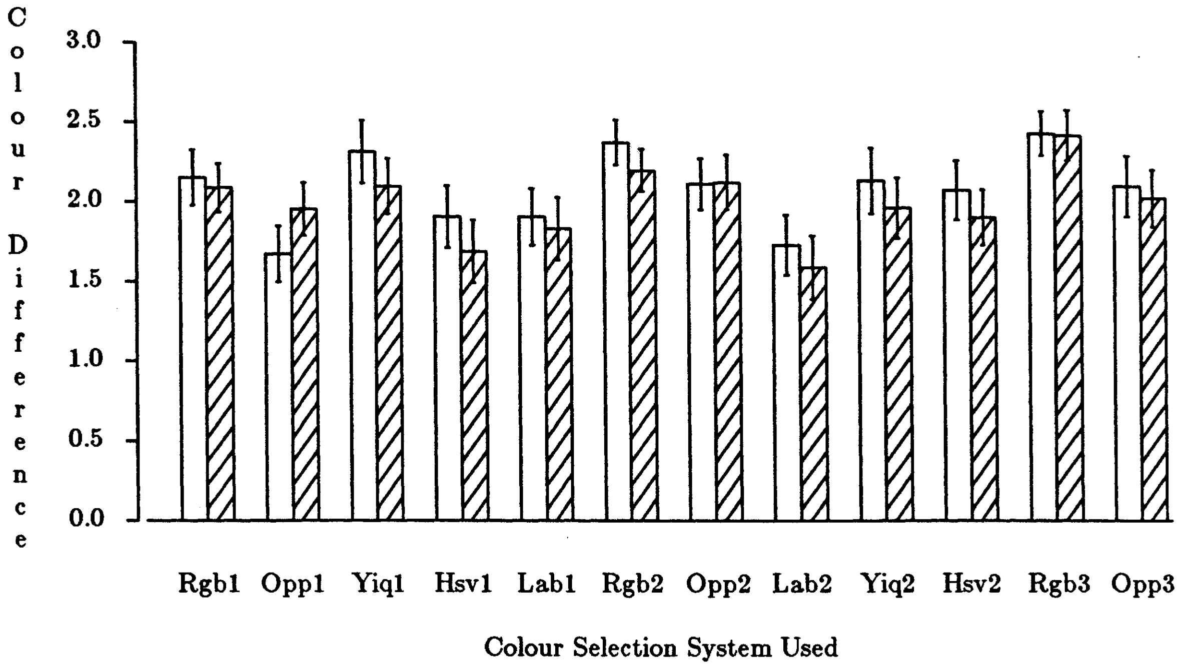

Figure 6.8b. Mean log colour difference at task completion by selection system and session half. The unshaded bars represent the means of those trials occuring in the first half of each session, while the shaded bars represent performance during the second half of each session. Incomplete trials are omitted.

As in Section 6.1.5, each figure in this section shows the proportion of trials for which a given threshold was reached and the average time taken by subjects to reach the threshold. In Figure 6.9a, these measures have been calculated across all trials by colour slice for a threshold value of 25 CIELAB colour difference units. Figure 6.9b shows the performance measures for the same threshold value but by colour selection system and session half. Trials where the white target colour was matched are not included in the calculations at this threshold value since the initial colour difference for these trials is less than 25 CIELAB colour difference units. In a similar manner, Figures 6.10a and 6.10b show the same performance measures for a colour difference threshold of 14 CIELAB units. Measures for a 6 CIELAB unit threshold are shown in Figures 6.11a and 6.11b. Trials where white was matched *are* included in the calculations for both these threshold values.

The data graphed in Figure 6.9a indicates that the mean time to reach a colour difference threshold of 25 CIELAB units decreases at a relatively steady rate across all colour slices. The proportion of trials that reach the threshold increases slightly between the first two colour slices but remains more or less constant thereafter. According to Figure 6.9b, it appears that there is very little difference in performance between session halves for all of the RGB and one of the Opponent Colours systems (i.e. "Opp1") at this threshold level. These systems ranked as the lowest 4 in terms of order of performance improvement between session halves. Subjects using the other selection systems seemed to improve their performance significantly between session halves. In spite of this apparent difference in learning between selection systems, the RGB and Opponent Colours based systems still ranked as 6 of the best 7 in the second session half according to the mean time taken to reach the 25 CIELAB unit threshold. However, the magnitude of the performance differences between the RGB and Opponent Colours systems and the other systems decreases considerably from the first session half to the second.

Figure 6.10a shows that performance, averaged across all selection systems, also improves steadily at the 14 CIELAB threshold level as subjects become more familiar with the selection systems. As in Figure 6.9a, the proportion measure in this figure also initially increases slightly and then levels off for the remainder of the session. One of the trends observed in Figure 6.9b can also be seen in Figure 6.10b: all the RGB systems and the "Opp1" Opponent Colours system again exhibit very little differences in performance between session halves. These systems rank as the lowest 4 in terms of improvement between session halves. Most of the other systems, on the other hand, showed considerable improvement.

At the 6 CIELAB level, the average performance measures calculated by colour slice are similar to those at previous levels: the mean time taken to reach the threshold decreases at a more or less uniform rate across all colour slices and the proportion of trials for which the threshold was reached rises slightly for the first two colour slices and then levels off (see Figure 6.11a). Figure 6.11b shows that the trends of previous thresholds are less clear at the 6 CIELAB unit level. However, the RGB and "Opp1" Opponent Colours systems still appear to exhibit slightly less improvement between session halves than the other systems. These systems rank as 4 of the lowest 5 in the order of improvement between session halves.

In general, it seems that the pattern of relative performance improvements (between colour selection systems) remained relatively constant for the different threshold levels. The overall magnitude of the performance improvements, however, seem to be somewhat less at the 6 CIELAB unit threshold than at the 14 and 25 CIELAB unit levels. These behavioral patterns are further discussed in Chapter 7.

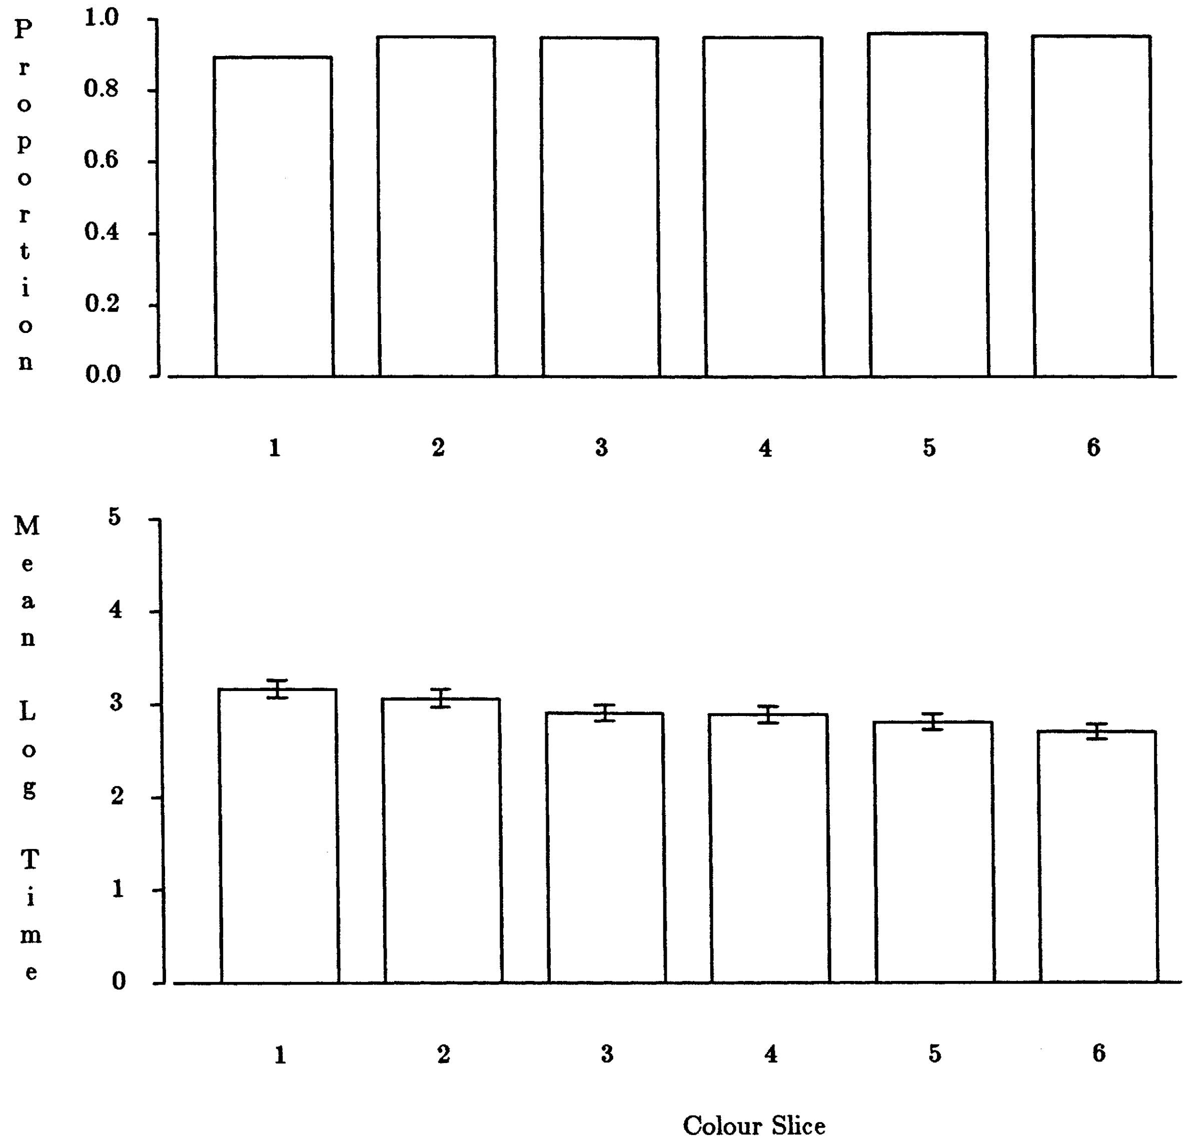

Figure 6.9a. Threshold of 25 LAB colour difference units by colour slice. Upper graph shows proportion of trials for which the threshold was reached. Lower graph shows mean log time taken to reach the threshold. Trials where white was matched were excluded from the calculation of these measures.

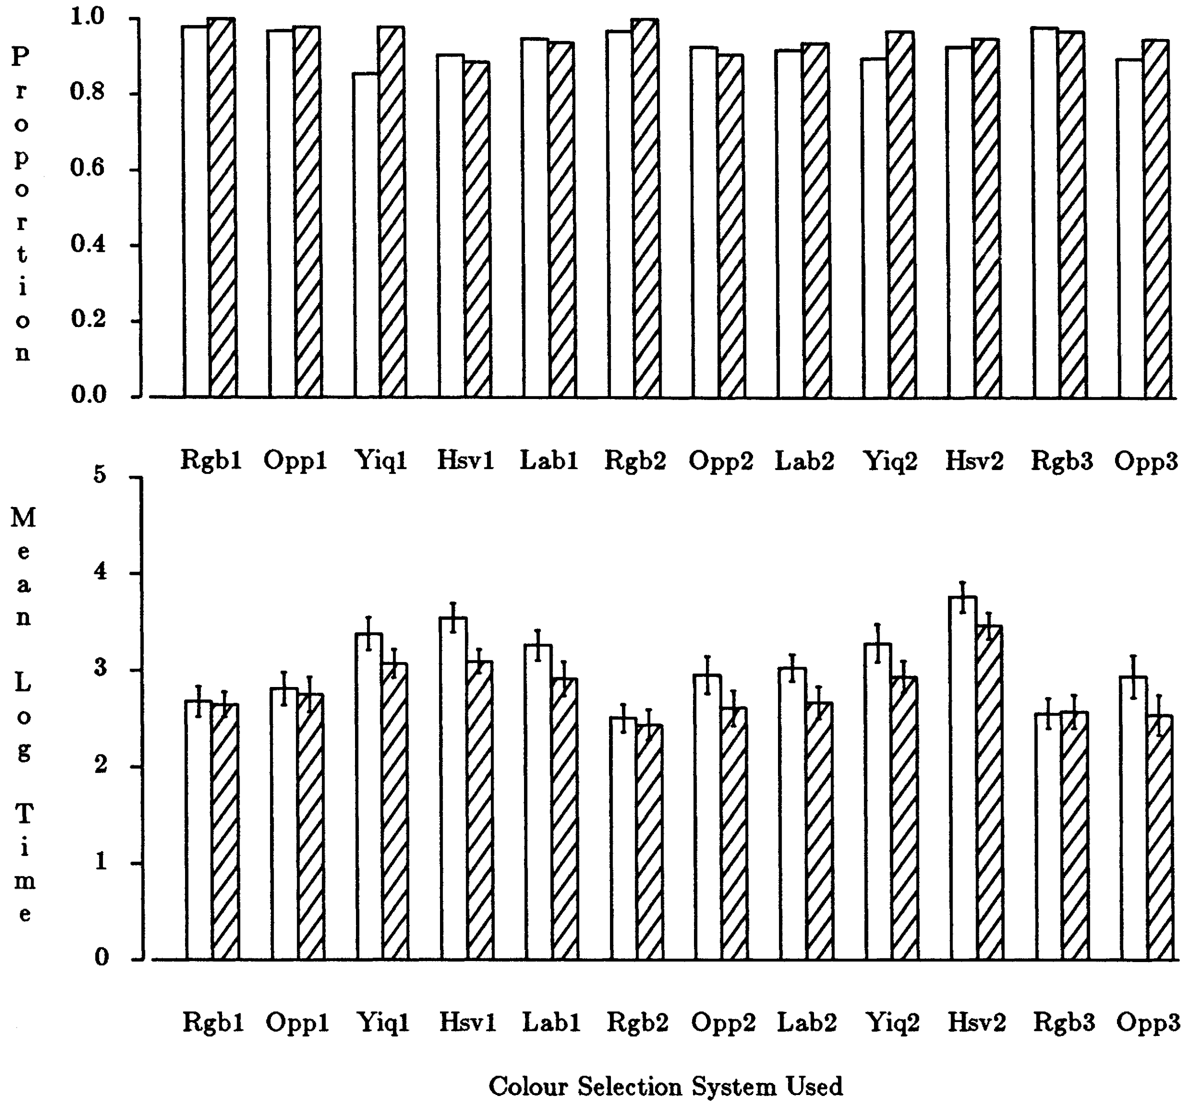

Figure 6.9b. Threshold of 25 LAB colour difference units by colour selection system and session half. Upper graph shows the proportion of trials for which the threshold was reached. Lower graph shows the mean log time taken to reach the threshold. The unshaded bars represent the measures for those trials occuring in the first half of each session, while the shaded bars represent performance during the second half of each session. Trials where white was matched are excluded from the calculations of these measures.

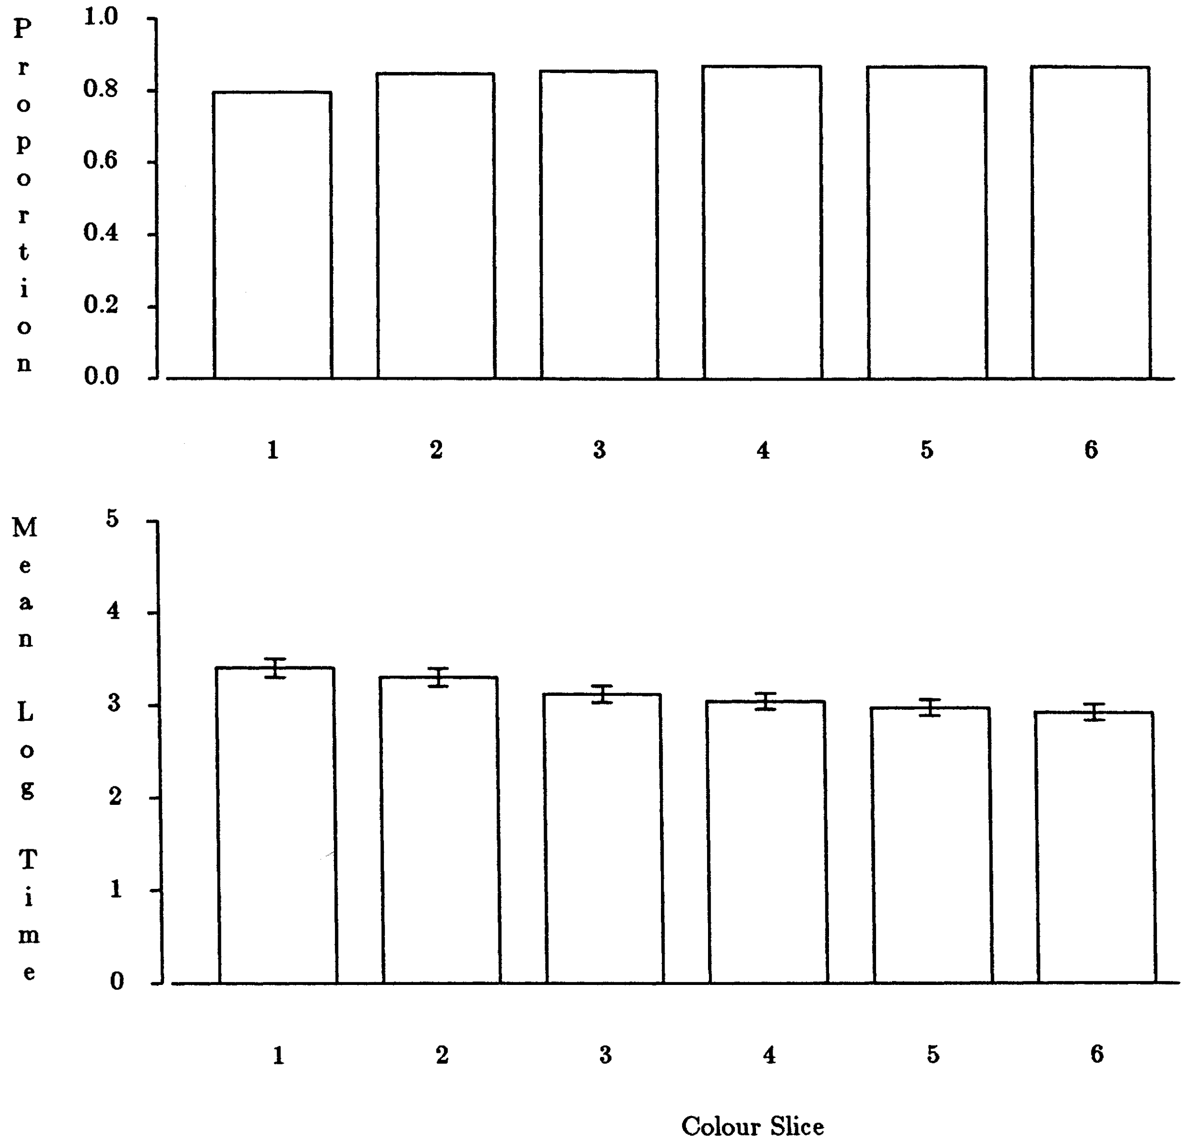

Figure 6.10a. Threshold of 14 LAB colour difference units by colour slice. Upper graph shows proportion of trials for which the threshold was reached. Lower graph shows mean log time taken to reach the threshold.

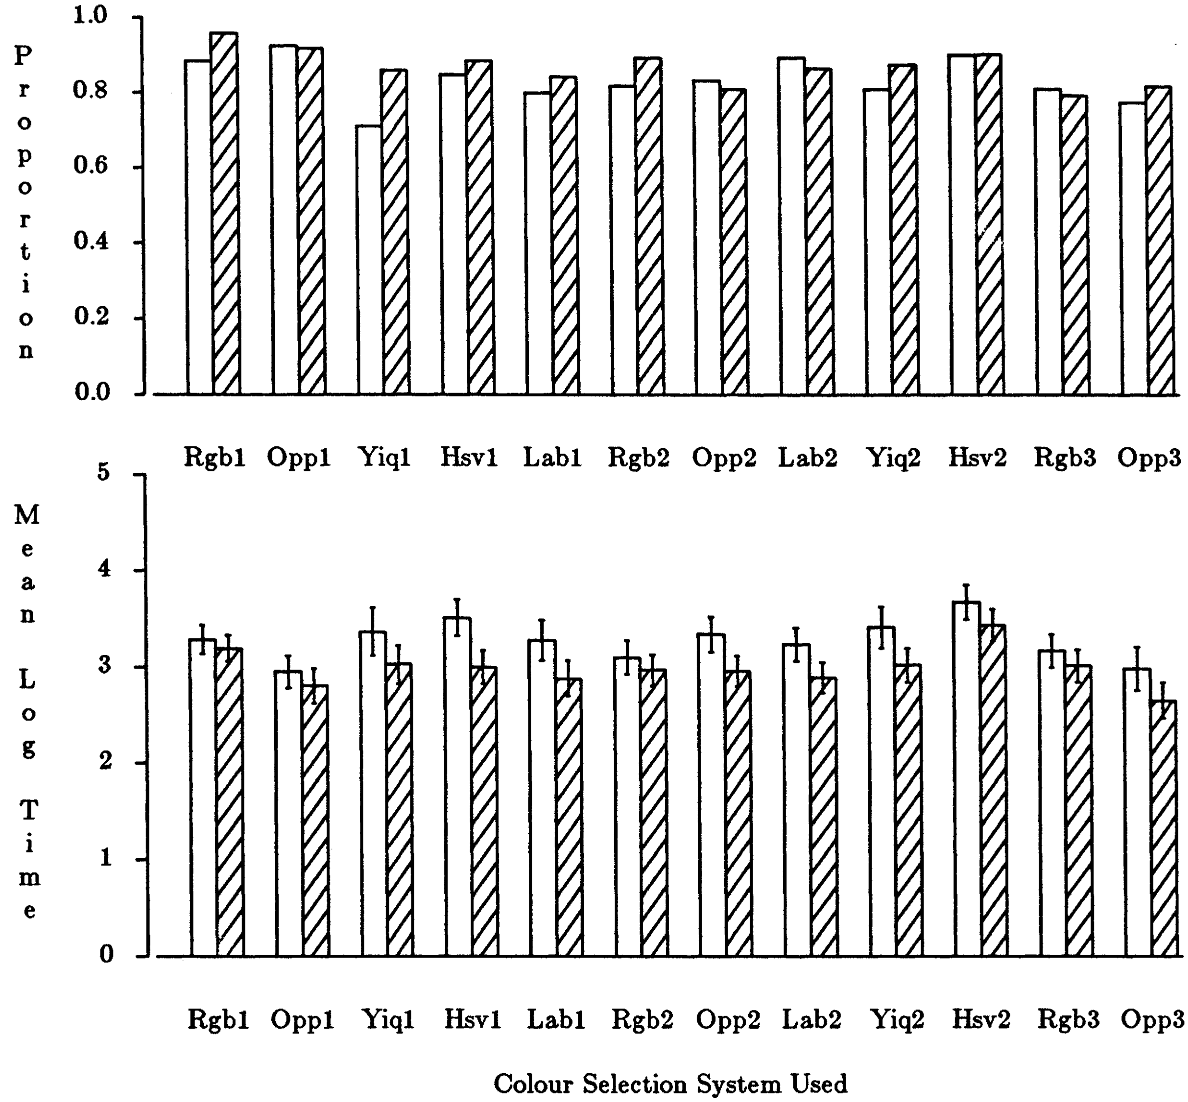

Figure 6.10b. Threshold of 14 LAB colour difference units by colour selection system and session half. Upper graph shows the proportion of trials for which the threshold was reached. Lower graph shows the mean log time taken to reach the threshold. The unshaded bars represent the measures for those trials occuring in the first half of each session, while the shaded bars represent performance during the second half of each session.

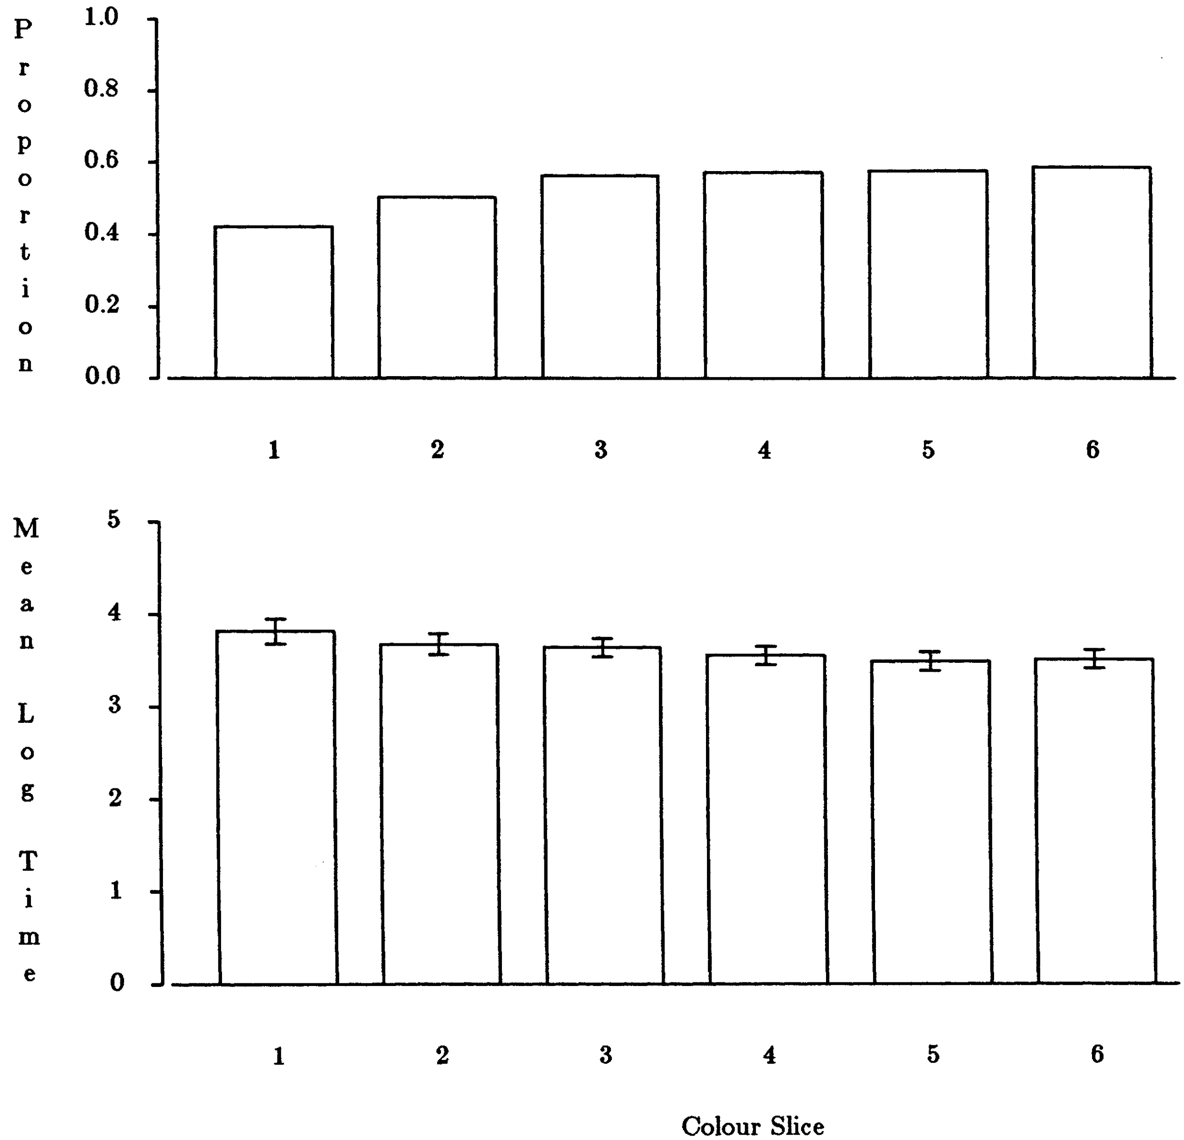

Figure 6.11a. Threshold of 6 LAB colour difference units by colour slice. Upper graph shows proportion of trials for which the threshold was reached. Lower graph shows mean log time taken to reach the threshold.

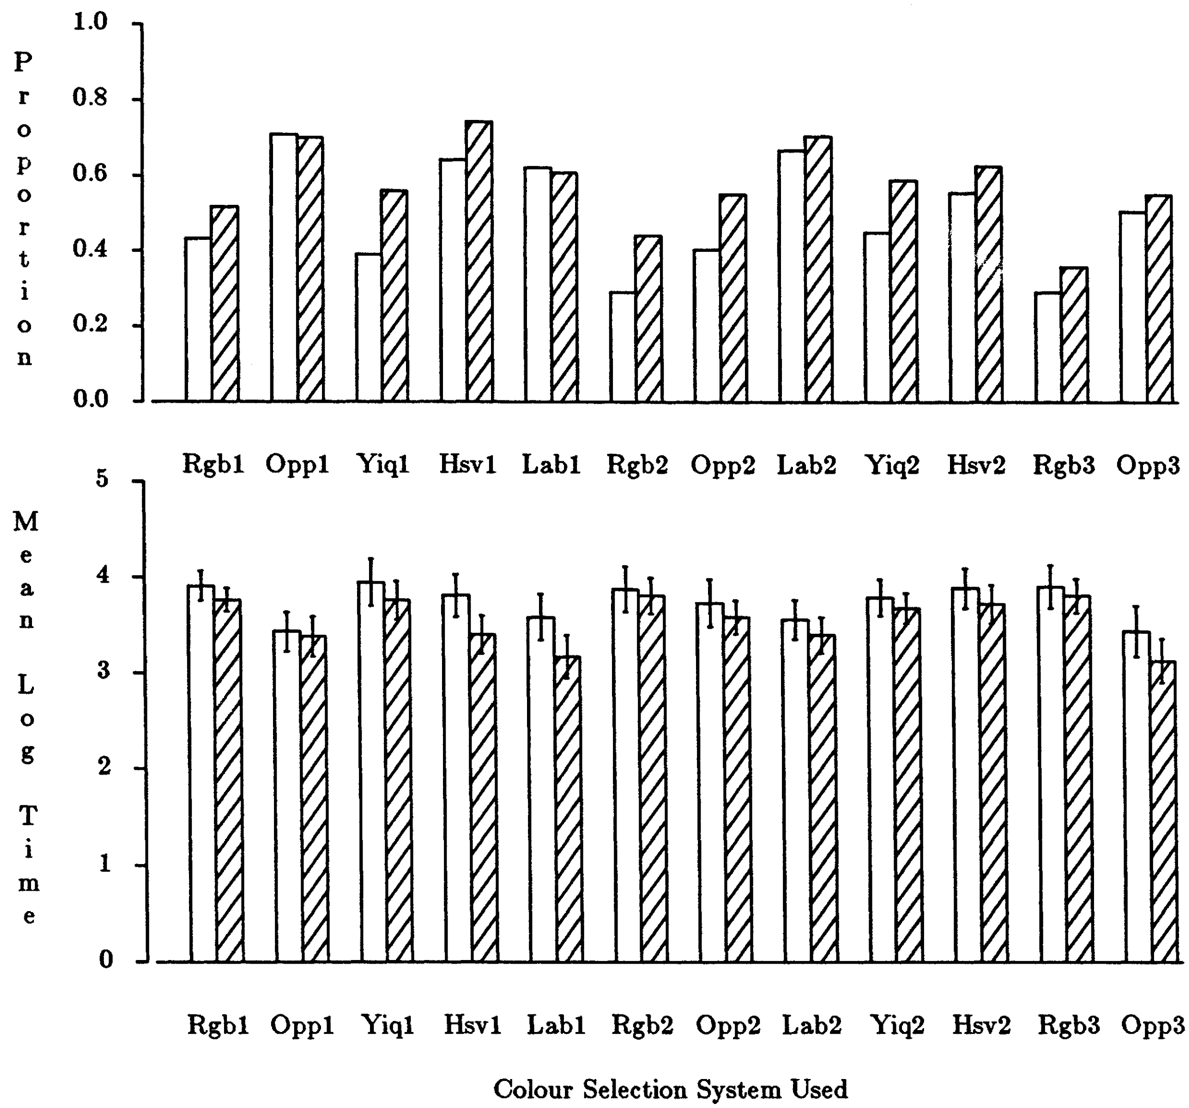

Figure 6.11b. Threshold of 6 LAB colour difference units by colour selection system and session half. Upper graph shows the proportion of trials for which the threshold was reached. Lower graph shows the mean log time taken to reach the threshold. The unshaded bars represent the measures for those trials occuring in the first half of each session, while the shaded bars represent performance during the second half of each session.

6.3. Summary of Comments

Each of the 96 subjects who participated in the experiment was asked to provide comments at the completion of the colour matching session. Figure 5.4 of the previous chapter showed the information form on which these comments were written. A summary of the comments made is given below in Table 6.3. The comments, as presented in the table, were combined into groups of similar meaning by the author. The reader should keep this in mind since this subjective interpretaion of the comments could be the source of some bias.

Summary of Comments	
Comment Description	Frequency
Compliments about the experiment.	12
Complaints of eye strain or that the experiment lasted too long.	9
It was difficult to make fine adjustments to the colour of the test stimulus.	9
The alphanumeric terminal or the colour moniter was too bright.	8
The colour properties subjects controlled to perform colour matching were difficult to comprehend or use.	5
Complaints about the manner in which the subject was informed that the 3 minute time limit had been exceeded.	5
The boundaries of the displayable colour gamut were frequently encountered.	5
Miscellaneous comments or suggestions about the tablet interface.	5
The tablet was difficult to use.	4
The feedback of the system in response to the movements of the puck seemed slow or inconsistent.	4
Miscellaneous comments on the difficulty of the assigned task.	4
The instructions were difficult to understand.	4
Complaints about the size or position of the test stimulus.	4
Miscellaneous complaints about various distractions that occurred during the experiment.	3
Certain colours were particularily difficult to match.	3
Complaints of hand or finger strain.	2
The level of light in the experiment room was too low.	2
Complaints about after-images.	2
Subject tended to use random movements of the puck in order to achieve a match.	2
Subject tended to manipulate only one degree of freedom at a time even though it was possible to alter two degrees of freedom simultaneously.	1
Subject tended to think in terms of a colour order system other than the one being used in the experiment.	1

Table 6.3. Summary of comments written by subjects at the completion of each session.

On the information form shown in Figure 5.4, subjects are asked specifically whether they had "any difficulty understanding the instructions." Hence, each of the 96 subjects was required to reply to this particular inquiry. From Table 6.3, it can be seen that a total of 4 subjects, or approximately 4% of the total sample, replied that there had been difficulty encountered in this regard. Since the proportion of subjects that had trouble understanding the instructions is quite low, it can therefore be inferred that the instuctions served adequately in preparing subjects for the colour matching tasks.

A number of subjects complained about the level of lighting during the experiment and several subjects mentioned that they had some trouble with eye strain. A correlation between these two phenomena is quite likely since the perceived brightness of CRT type display devices seems to be more difficult to control at lower intensity levels. The visual fatigue problems encountered are not thought to have greatly affected the results of the experiment since the effects of these difficulties were most likely concentrated near the end of each session and can therefore be isolated. The relationship between low ambient light levels and eye strain is also discussed in the next chapter.

6.4. Performance of Experienced Subjects

In this section, the performance of the 2 subjects from the "experienced" category is analyzed.

One subject, who will be referred to as Subject 1 in the remainder of this work, has been involved in colour matching experiments such as this one for approximately 30 years. In fact, much of this work has involved participation in experiments where it was specifically required to make "optimal" colour matches. It is important to note that all this experience in performing colour matching tasks was based on the use of the RGB colour model. The subject did not have any experience using selection systems based on other colour models. The subject did, however, have a fundamental understanding of these other models. This subject had never used a graphical input device, such as a tablet or mouse, before participating in this experiment.

The other subject, subsequently referred to as Subject 2, has spent many years using tablets and mice to perform a variety of interactive tasks. Compared to Subject 1, however, this subject had little experience in performing colour matching prior to participating in this experiment.

Both of the experienced subjects had a fundamental understanding of the concepts underlying each of the colour order systems and interaction techniques employed in the experiment. Furthermore, both subjects were made aware of which model and interaction technique were to be used before the start of each session. Subject 1 performed the experiment using 9 of the 12 colour selection systems over a period of two days. Subject 2 performed the colour matching tasks using 10 of the selection systems. These latter sessions were performed one or two per day but spread out over a period of several months.

6.4.1. Task Incompletions

Table 6.3 shows, by colour selection system, the number of trials where the experienced subjects failed to complete the colour matching task within the allowed 3 minute time limit. From this data it can be seen that Subject 2 timed out considerably more often than Subject 1. In fact, Subject 1 only failed to complete the required task in the session where the "single attribute control" version of the YIQ system (i.e. "Yiq1") was used. Subject 2 tended to time out more when using selection systems based on the "multiple attribute control" interaction technique.

Task Incompletions		
System	Subject 1	Subject 2
Rgb1	0	1
Opp1	0	0
Yiq1	2	1
Hsv1	0	1
Lab1	0	2
Rgb2	0	5
Opp2	0	0
Lab2	0	6
Yiq2	0	0
Hsv2	-	3
Rgb3	-	-
Opp3	-	-
Total	2	19

Table 6.3. Number of times experienced subjects exceeded the allowed time limit by colour selection system.

6.4.2. Trial Duration

Figures 6.12a and 6.12b show, by colour selection system, the mean times taken to complete the colour matching tasks for Subject 1 and Subject 2 respectively. Time is measured in seconds. Incomplete trials are not included in the calculation of the means.

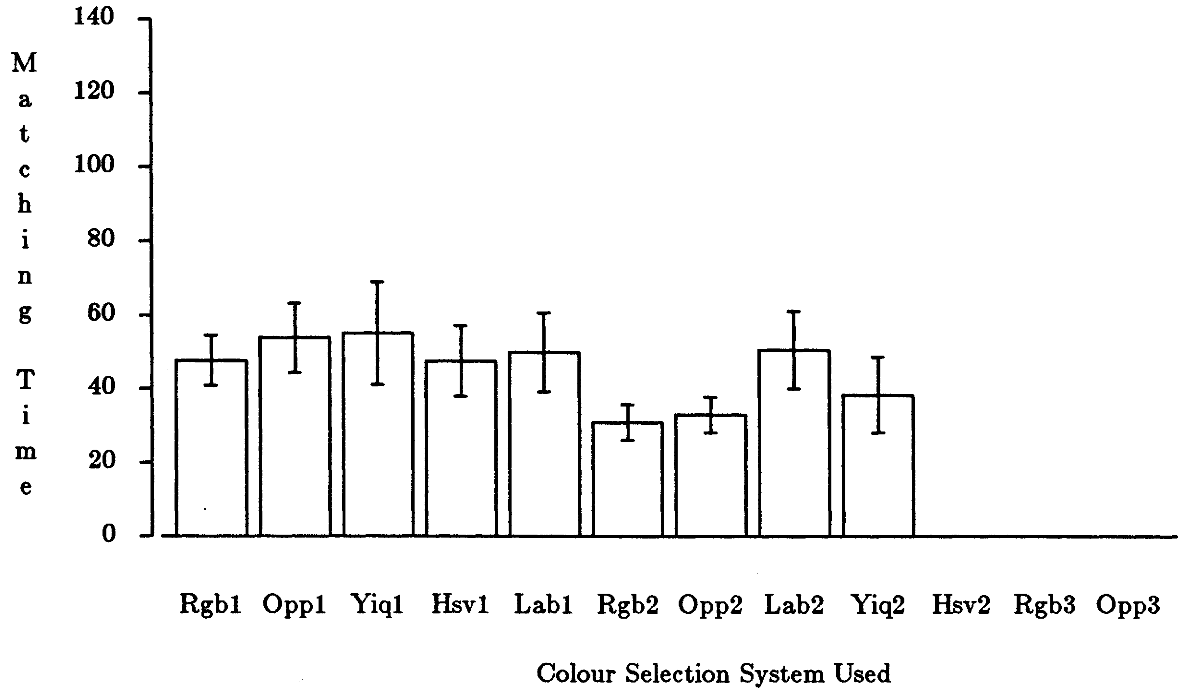

Figure 6.12a. Mean time taken to complete colour matching task by colour selection system for experienced Subject #1. Incomplete trials are omitted.

The most obvious effect that can be seen in this data is the difference between subjects in the time spent performing the colour matching tasks. Subject 2 used considerably more time, on average, to complete the tasks than Subject 1. This trend holds true across almost all the selction systems.

The differences in mean trial duration between selection systems are rather insignificant for Subject 1. However, performance using the "multiple attribute control" versions of the RGB and Opponent Colours systems seemed to be slightly better than for the other systems. These systems ranked as the best 2 for the criteria of minimizing mean trial duration. The "single attribute" versions of RGB and Opponent Colours ranked 5th and 8th respectively. For Subject 2, the differences in mean trial duration between systems are larger but not necessarily more significant. The best 3 systems for Subject 2, according to this criteria, are based on the RGB and Opponent Colours colour order systems.

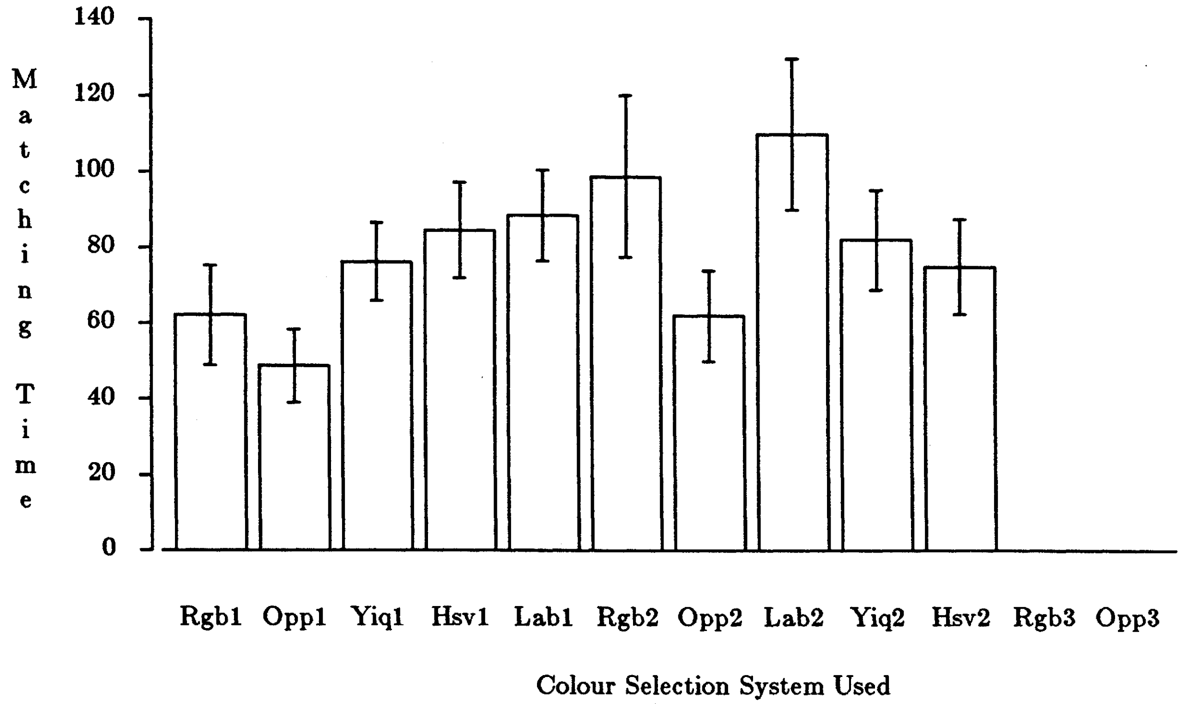

Figure 6.12b. Mean time taken to complete colour matching task by colour selection system for experienced Subject #2. Incomplete trials are omitted.

6.4.3. Colour Difference

The mean colour difference at trial completion for Subject 1 and Subject 2 are displayed by selection system in Figures 6.13a and 6.13b. Again, a logarithmic transformation has been applied to the data to achieve a normal distribution. Trials where the subjects did not complete the task in the allowed time were not included in the calculation of these performance measures.

In comparing the data of the two experienced subjects, there are no major differences in performance between subjects according to this measure. However, Subject 1 did seem to match significantly closer than Subject 2 using the "Rgb1" and "Opp1" systems.

Subject 1 appeared to perform significantly better in terms of colour difference using the selection systems based on RGB or Opponent Colours (these systems ranked as the best 4) than when performing the colour matching task using systems based on the other colour models. There are no major performance differences between selection systems in the data for Subject 2.

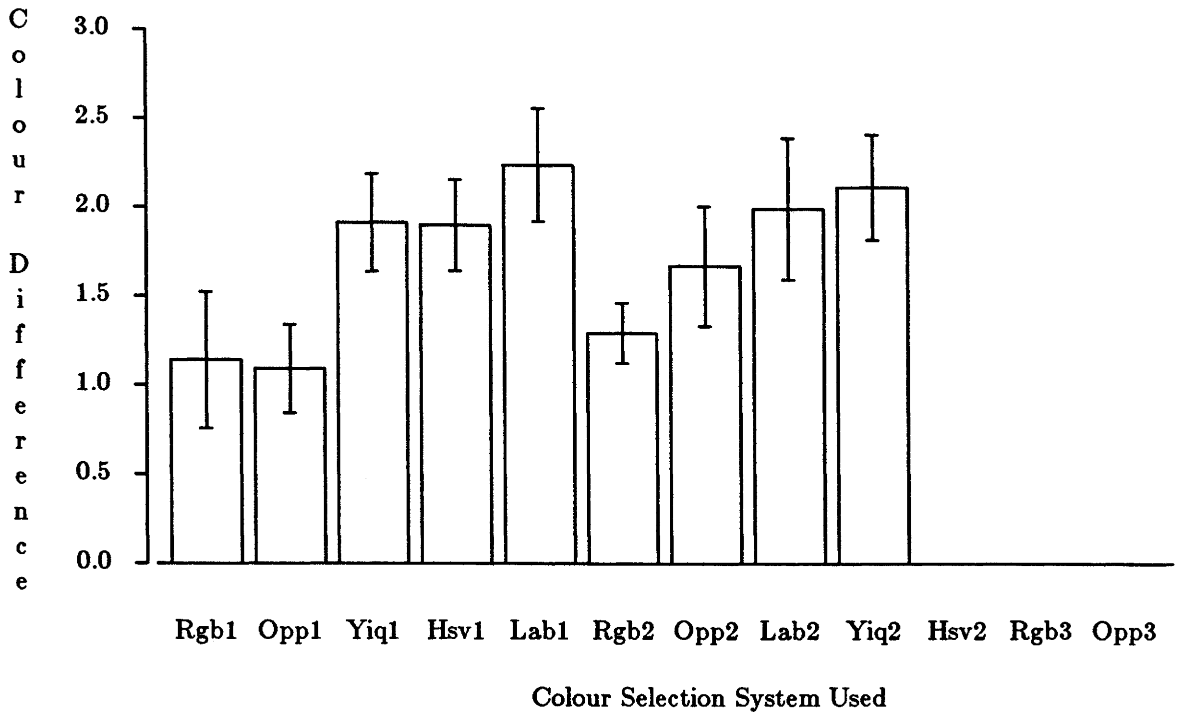

Figure 6.13a. Mean log colour difference by selection system for experienced Subject #1. Incomplete trials are omitted.

6.4.4. Colour Difference Thresholds

Figures 6.14a and 6.14b show the proportion of trials for which a threshold of 25 CIELAB colour difference units was reached and the average time taken by the subject to reach this threshold value. Figure 6.14a displays these values for Subject 1, while Figure 6.14b plots the values for Subject 2. Similarily, these performance measures are displayed for a threshold value of 14 CIELAB units in Figures 6.15a and 6.15b. Figures 6.16a and 6.16b show the same measures for a colour difference threshold of 6 CIELAB units. Trials where the subjects failed to complete the colour matching task in the allowed time *are* included in the calculation of the performance measures presented here. A logarithmic transformation was applied to the times taken to reach a colour difference threshold in order that a more normally distributed data sample could be achieved.

At the 25 CIELAB unit threshold level, there were no major differences in performance between subjects for either performance measure. Of particular interest is the fact that Subject 2 tended to reach the threshold using the RGB based systems in the same time that it took Subject 1 to do the same (in spite of Subject 1's experience in using RGB systems).

Both subjects tended to reach the 25 CIELAB unit threshold faster using the RGB and Opponent Colours based selection systems than when using the other systems. In particular, these systems ranked as the best 4 for both subjects. However, the RGB systems seemed to

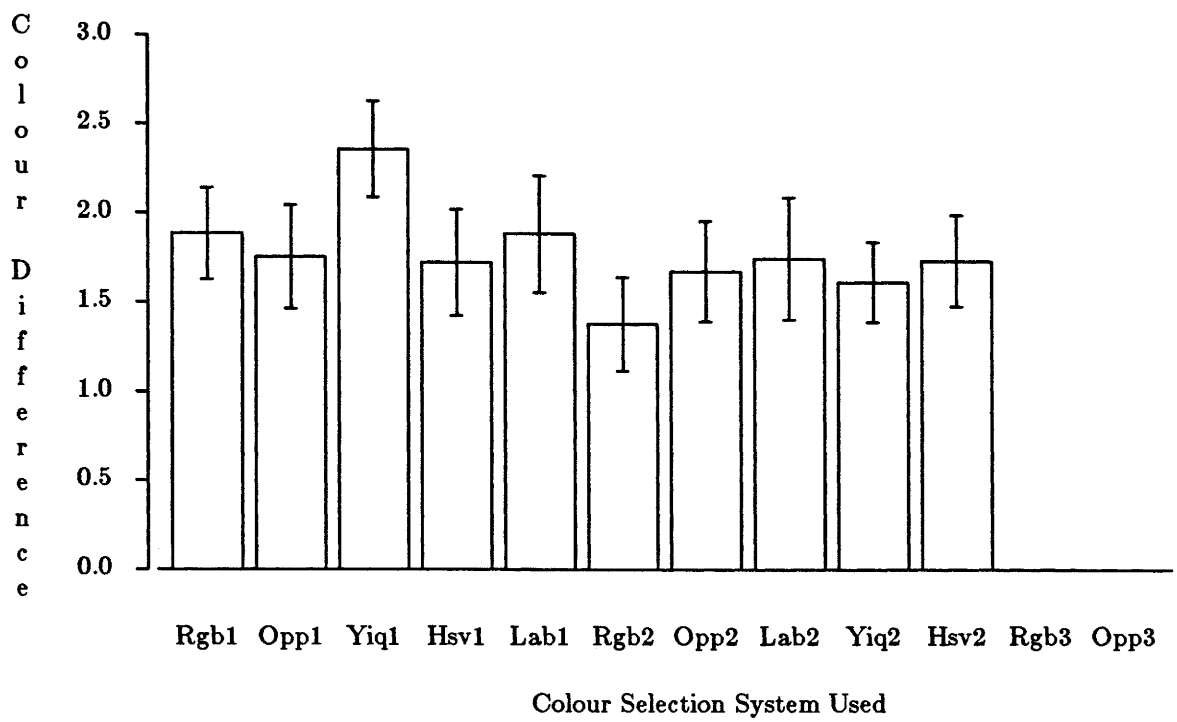

Figure 6.13b. Mean log colour difference by selection system for experienced Subject #2. Incomplete trials are omitted.

stand out more for Subject 1 when compared to this subject's performance using the other systems.

At the 14 CIELAB unit threshold, performance for the experienced subjects was very similar to that at the 25 CIELAB level. Performance between subjects tended to be more or less the same for all selection systems including those based on RGB. The RGB and Opponent Colours systems ranked as 4 of the best 5 for both subjects. Subject 1 still seemed to do especially well using the RGB based systems but the differences in performance between these systems and the others were noticably smaller at this thresold level.

At 6 CIELAB units, Subject 1 tended to reach the threshold faster than Subject 2 for a number of systems including both RGB based systems. For Subject 1, only the RGB systems appeared to stand out in terms of mean time to reach the colour difference threshold. The proportion measure at this threshold level also seems to indicate that Subject 1 performed moderately better using the RGB based systems than when using systems based on other colour models. There did not appear to be any significant differences in performance between selection systems for Subject 2 at the 6 CIELAB unit level.

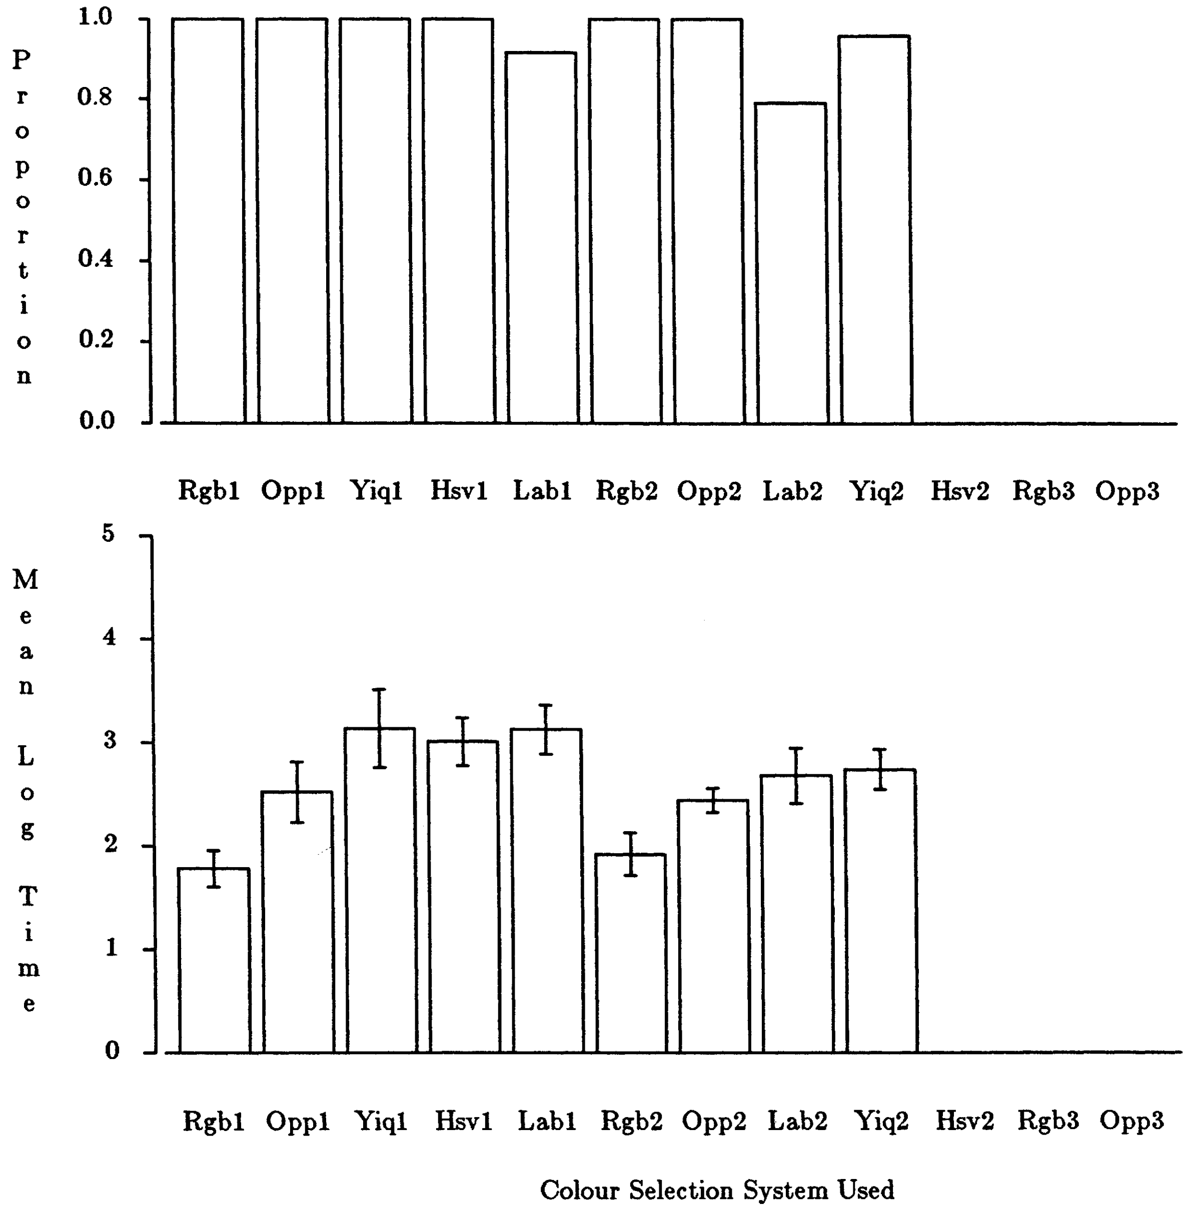

Figure 6.14a. Threshold of 25 CIELAB colour difference units by colour selection system for experienced subject #1. Upper graph shows proportion of trials in which the threshold was reached. Lower graph shows mean log time taken to reach the threshold. Trials where white was matched are excluded from these measures.

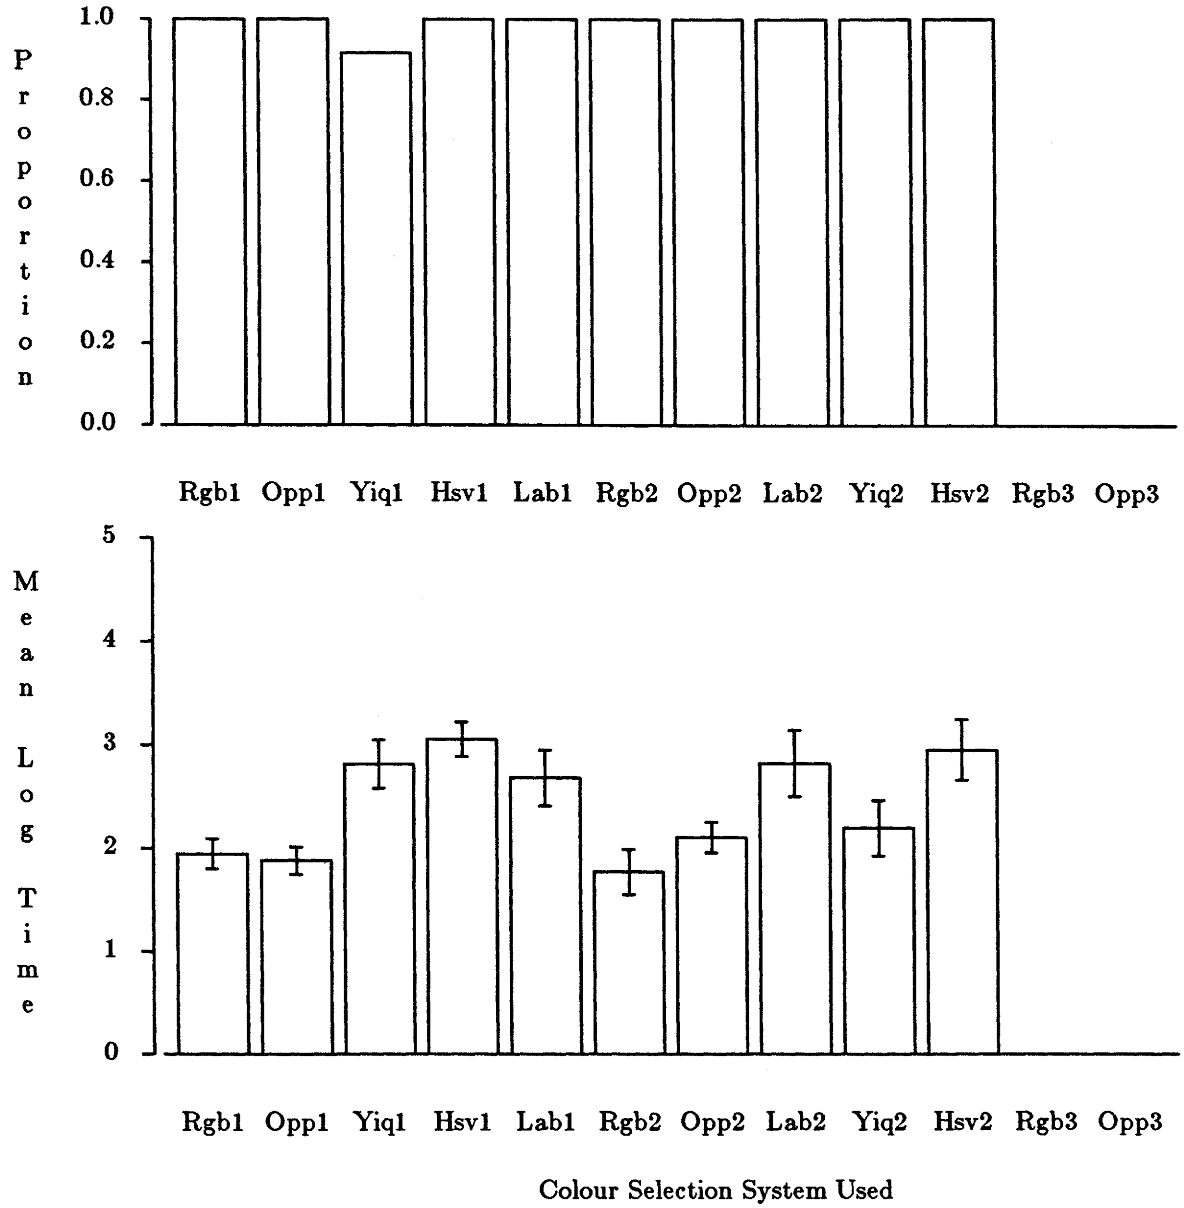

Figure 6.14b. Threshold of 25 CIELAB colour difference units by colour selection system for experienced subject #2. Upper graph shows proportion of trials in which the threshold was reached. Lower graph shows mean log time taken to reach the threshold. Trials where white was matched are excluded from these measures.

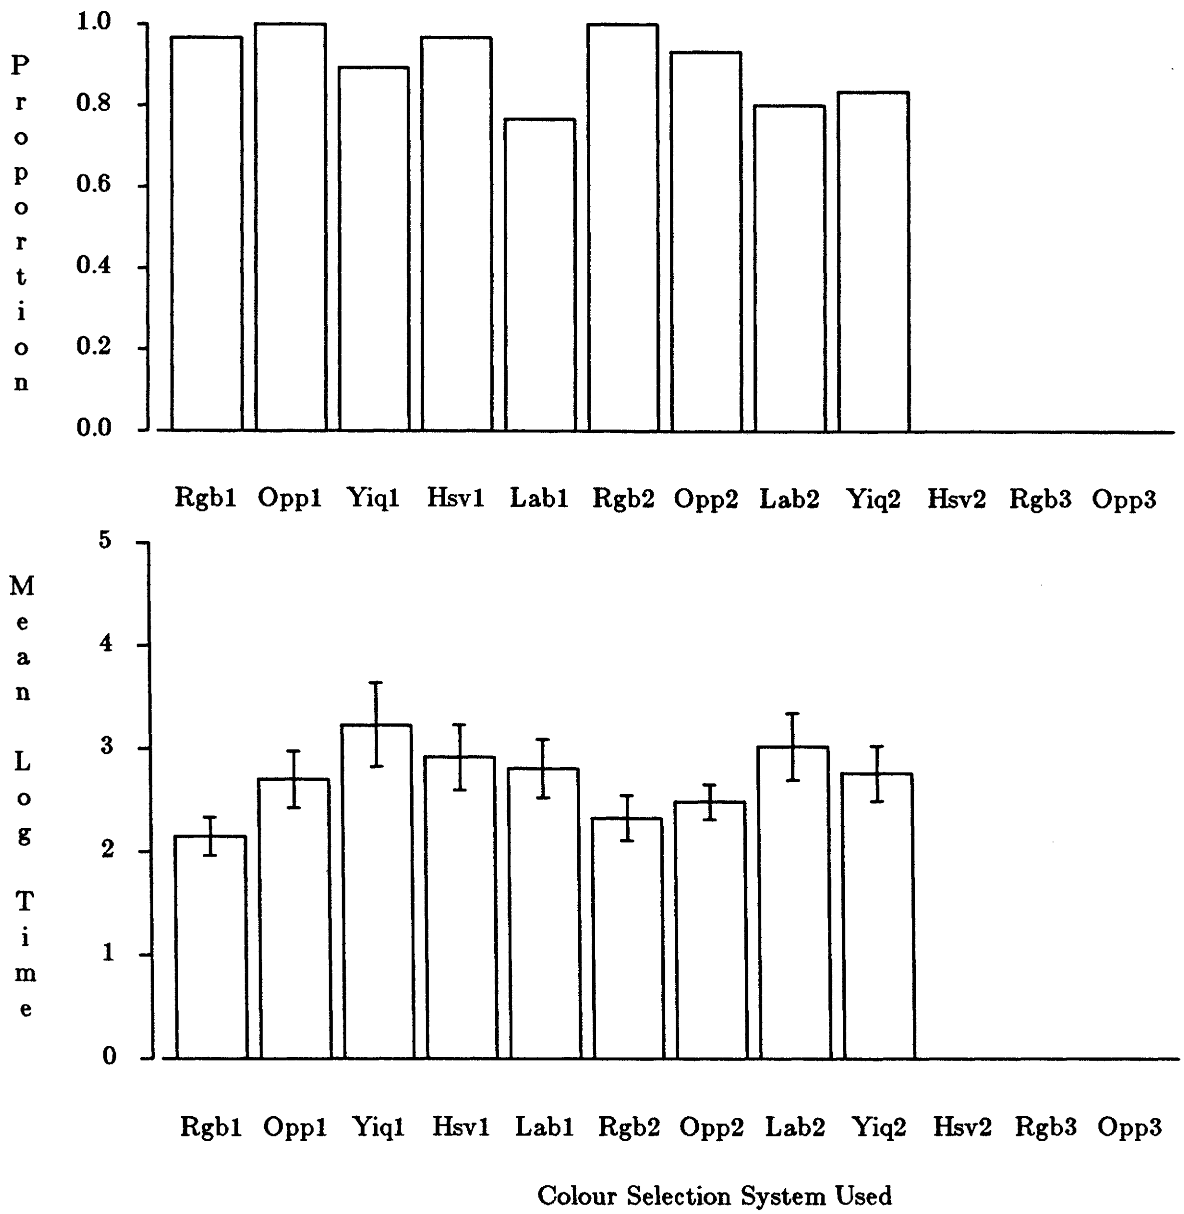

Figure 6.15a. Threshold of 14 CIELAB colour difference units by colour selection system for experienced subject #1. Upper graph shows proportion of trials in which the threshold was reached. Lower graph shows mean log time taken to reach the threshold. Trials where white was matched *are included* in the calculation of these performance measures.

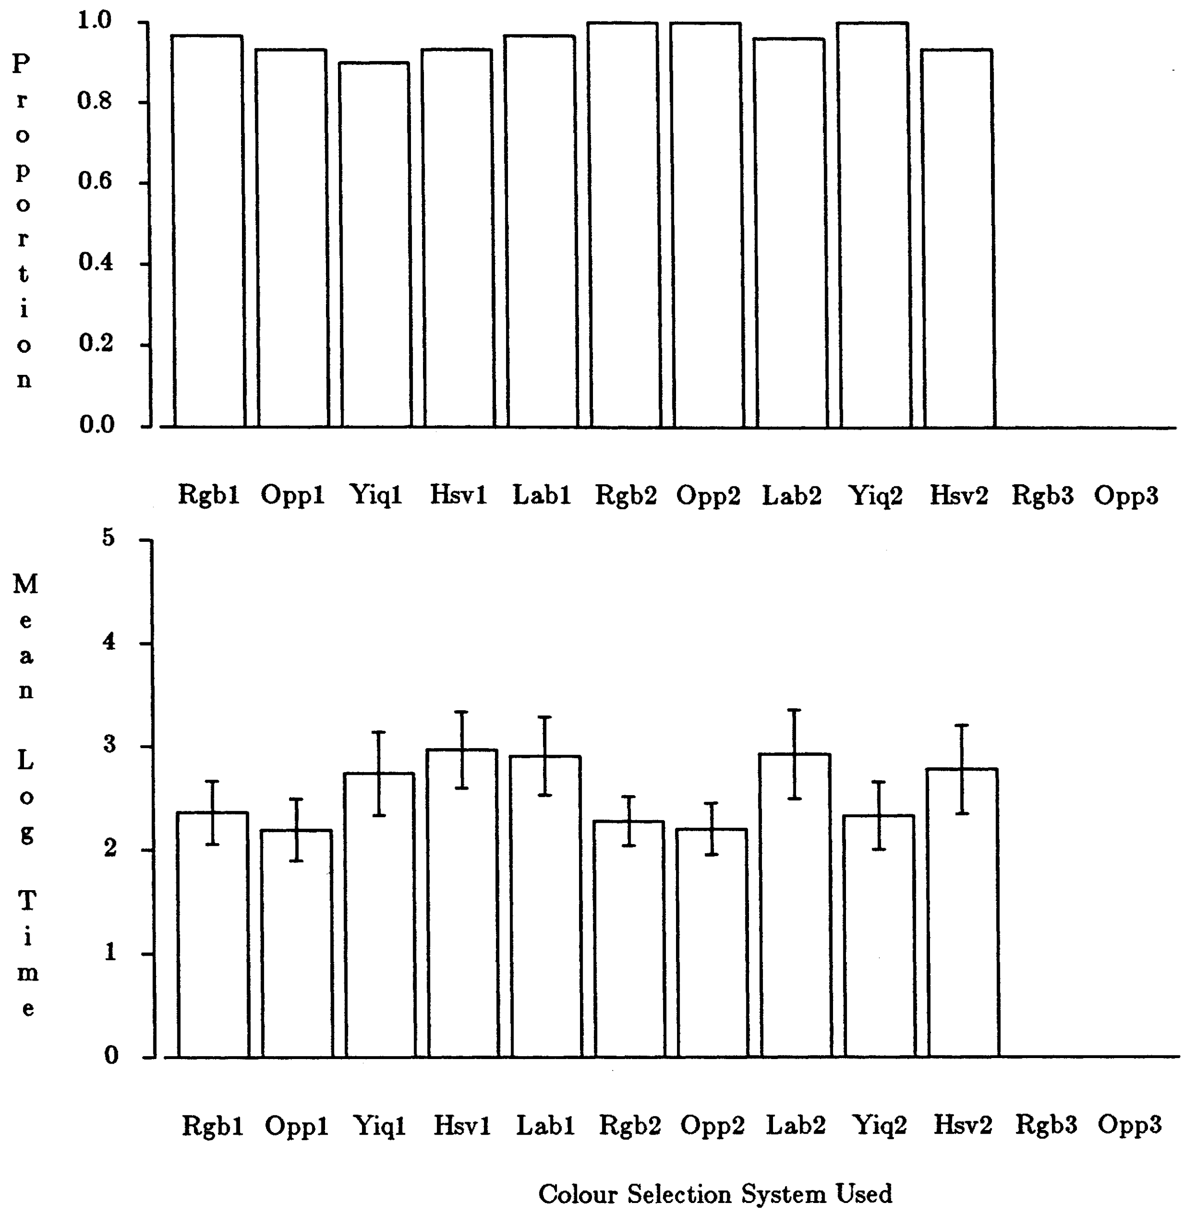

Figure 6.15b. Threshold of 14 CIELAB colour difference units by colour selection system for experienced subject #2. Upper graph shows proportion of trials in which the threshold was reached. Lower graph shows mean log time taken to reach the threshold. Trials where white was matched *are included* in the calculations of these performance measures.

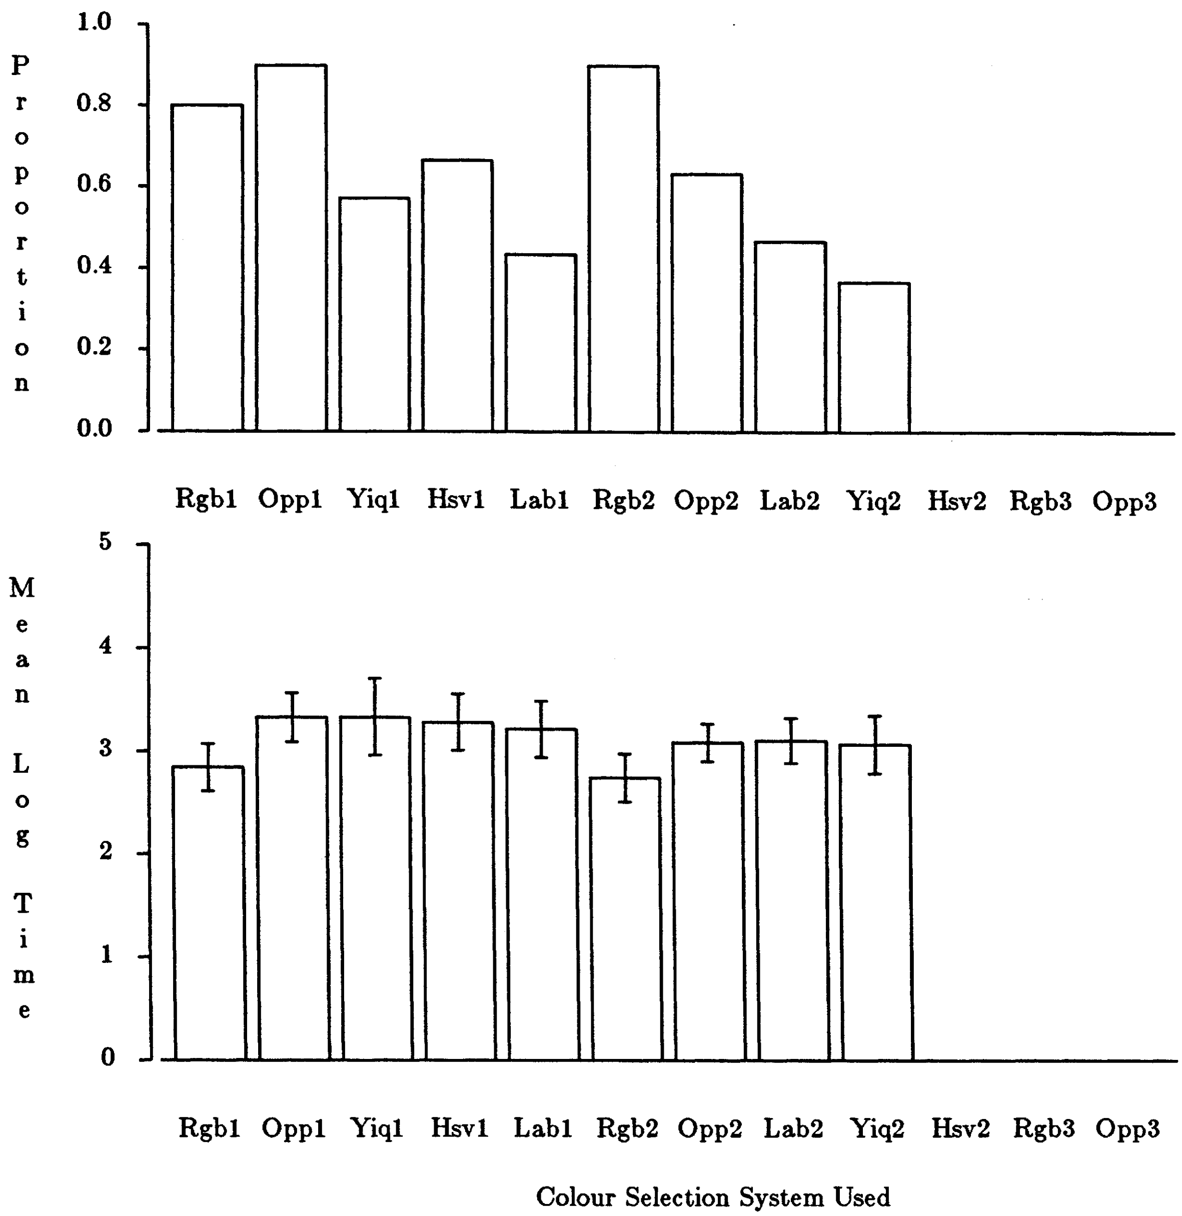

Figure 6.16a. Threshold of 6 CIELAB colour difference units by colour selection system for experienced subject #1. Upper graph shows proportion of trials in which the threshold was reached. Lower graph shows mean log time taken to reach the threshold. Trials where white was matched *are included* in the calculations of these performance measures.

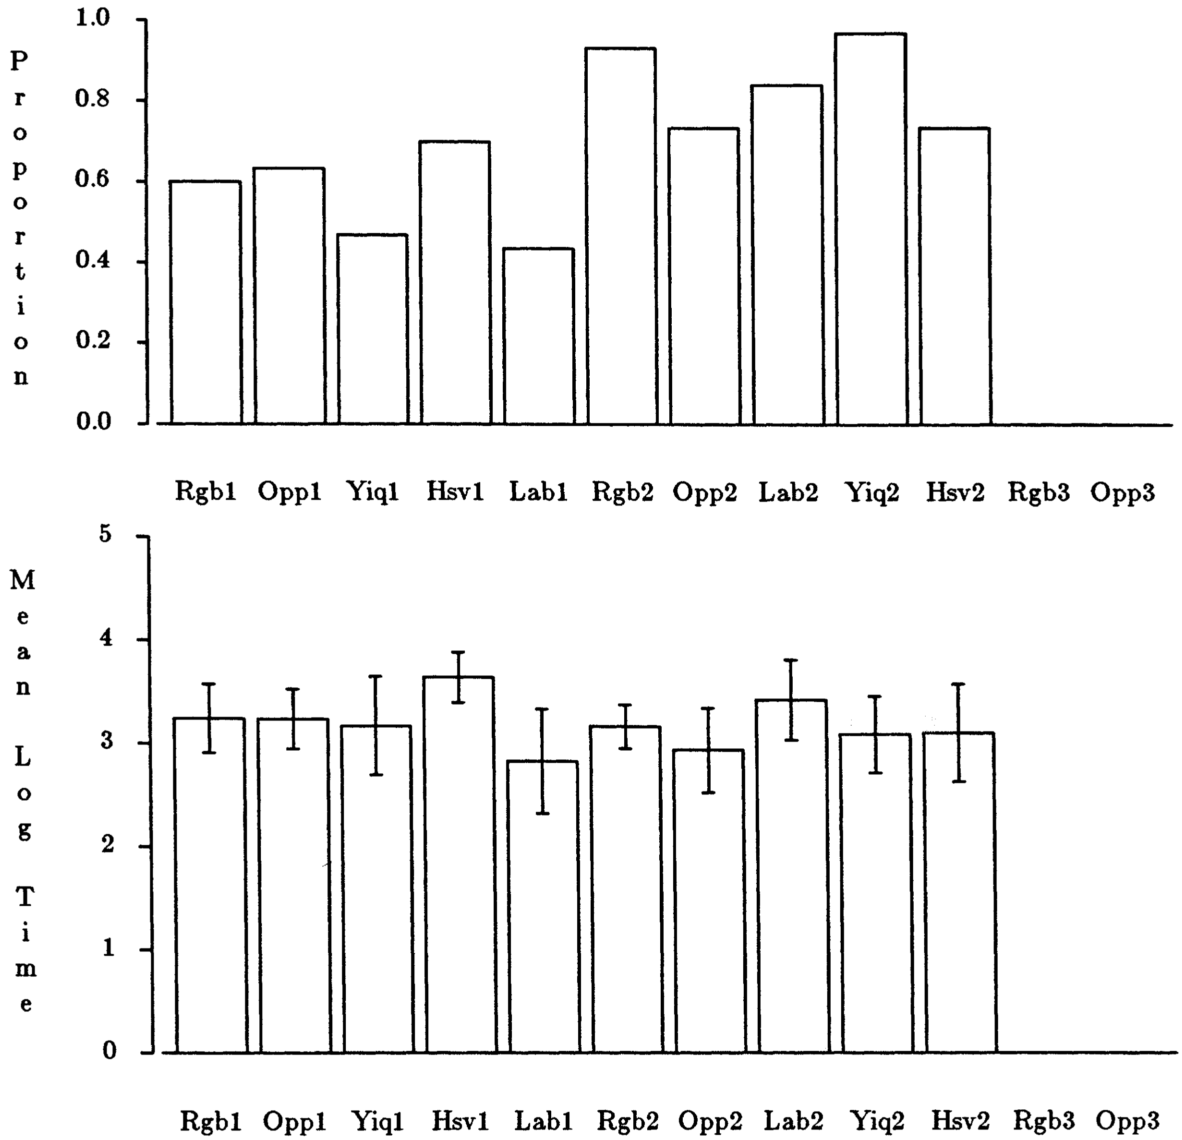

Figure 6.16b. Threshold of 6 CIELAB colour difference units by colour selection system for experienced subject #2. Upper graph shows proportion of trials in which the threshold was reached. Lower graph shows mean log time taken to reach the threshold. Trials where white was matched ***are included*** in the calculations of these performance measures.

7. Discussion

The preceding chapter presented a number of performance measures that quantified, according to several criteria, the behavior of subjects that performed the colour matching experiment. In this chapter, several conclusions about the observed behavior and the implications of this behavior on the user interface issues affecting interactive colour selection are derived from the results.

The first section of this chapter contains conclusions which can be inferred directly from the data presented in the previous chapter. In the second section, a model is composed that attempts to characterize the behavior of subjects during the experiment. A number of conclusive statements based on this model are then presented. Section three consists of a speculative discussion that seeks to explain some of the observed behavior in terms of what is known about the human visual system and other existing theories of human behavior.

7.1. Immediate Conclusions

The conclusions discussed in this section were deduced directly from the results. All inferences about how well subjects performed the colour matching task are based solely on the statistical measures presented in Chapter 6.

7.1.1. Separate Degree of Freedom Permits Finer Control of Lightness

In Section 6.1.4, the mean final colour difference measure was split into the three components of "hue," "chroma" and "lightness" (as defined by the standard CIELAB colour difference formulas [CIE78]). A number of very clear effects can be seen in this data that allow conclusions to be made about which properties of colour it is useful to have users control via a separate degree of freedom. It should be noted that this decomposition of the final colour difference was only performed for the data of the inexperienced subjects. Unless otherwise stated then, the conclusions presented below that are based on these measures should be interpreted to apply only to the performance of inexperienced subjects since experienced users may very well exhibit completely different performance characteristics. In addition, the reader should remember that the attributes of hue and chroma for which differences were measured are substantially different from the colour properties subjects controlled when using selection systems based on the CIELAB colour model (see Section 6.1.4 for a description of the hue, chroma and lightness difference measures).

The most prominent effect in the data presented in Section 6.1.4 appears in the "lightness" component of the mean final colour difference. Here it can be seen that, in general, subjects who performed the experiment using colour selection systems that had an explicit control for a colour property similar to CIELAB "lightness" matched significantly closer in terms of lightness than subjects who used a system that did not have a specific control for this attribute. In this experiment all the colour selection systems had such a separate degree of freedom for "lightness," "luminance" or "value" except the systems based on the RGB colour model.

Since the colour difference at the completion of the colour matching task can be regarded as a measure of how easy it was for subjects to make small refinements to the test stimulus, it can be concluded that the subjects in this experiment were better able to "fine tune" colours in terms of the lightness attribute if the colour selection system they used had an explicit control for a lightness-like property.

The implication of this result for the design of human-computer interfaces is fairly obvious. For colour selection tasks where it is important to have fine control over the *intensities* of the colours used in an image, the colour selection system implemented should have a separate degree of freedom for an "intensity" property. Although there is no evidence here to support it, one would intuitively expect the above statement to apply for experienced users as well.

7.1.2. Separate Degrees of Freedom Do Not Imply Finer Control for Hue or Chroma

Given the previous argument that direct control of "lightness" allowed subjects to match colours closer in terms of CIELAB lightness, it is not unreasonable to expect that separate degrees of freedom for other colour properties would also permit subjects to achieve relatively closer matches in terms of these attributes. For the CIELAB defined properties of "hue" and "chroma," however, this appears *not* to be the case.

The colour selection systems based on the HSV colour model have separate degrees of freedom for properties (i.e. hue and saturation) that are quite similar to the CIELAB versions of "hue" and "chroma." According to the data presented in Section 6.1.4, subjects using these selection systems did not match significantly closer in terms of hue and chroma than subjects who used selection systems that did not allow direct control over these colour properties. In fact, none of the selection systems investigated in this experiment seemed to really stand out for the criteria of minimizing hue and chroma differences.

Hence, none of these selection systems, particularly the systems with separate degrees of freedom for "hue" and chroma", allowed subjects to exercise finer control over the hue and chroma properties of colour.

For colour selection tasks that require fine control over chromatic properties such as CIELAB "hue" and "chroma," it does not appear to be advantageous to give users explicit controls for these attributes. In addition, there is evidence to believe that the colour model on which the selection interface is based does not appear to have a significant effect on how easily users can make colour refinements in terms of hue and chroma.

7.1.3. RGB and Opponent Systems Have Similar Desirable Performance Properties

One pattern that seems to be repeated in the results for several performance criteria is that subjects who used colour selection systems based on the RGB or Opponent Colours colour models tended to perform "better" than subjects that used selection systems based on other models. Of the measures applied to the data of the inexperienced subjects in Section 6.1, the RGB and Opponent Colours based systems seemed to stand out among the "best" for: the number of task incompletions, the mean trial duration times and the mean colour difference thresholds for 25 and 14 CIELAB units. The data shown in Section 6.2 also seems to indicate that, for subjects who used the RGB and Opponent Colours systems to perform the colour matching tasks, there were relatively small differences in performance between the first and second halves of a session according to a number of criteria. In particular, the mean final colour difference and the mean time taken to reach the colour difference thresholds of 25, 14 and 6 CIELAB units tended to remain relatively constant between session halves for these selection systems. Furthermore, these patterns in task performance are more or less repeated in the data for experienced subjects presented in Section 6.4.

Thus, it appears that the selection systems based on the RGB or Opponent Colours colour order systems have a number of similar desirable performance characteristics. It should be noted, however, that there are a number of criteria for which these systems did not rank so well. Of particular interest is the observation that, although subjects using the RGB or Opponent Colours systems were able to reach the colour difference thresholds of 25 and 14 CIELAB units faster than subjects using other selection systems, they were not able to do so for the 6 CIELAB unit threshold. A discussion of why these systems seemed to perform well for some criteria and poorly for others is deferred until the next section where a model of subject behavior is developed to help understand these performance trends.

7.1.4. HSV, YIQ and CIELAB Systems Require Learning to Use Effectively

As was previously noted, subjects using RGB or Opponent Colours based systems generally showed relatively small differences in performance between the first and second session halves according to several of the performance measures discussed in Section 6.2. Conversely, for subjects who used systems based on the other colour models (i.e. HSV, YIQ and CIELAB), there was a significant improvement in performance between session halves according to the same criteria. For several of these measures, particularly the colour difference thresholds of 25 and 14 CIELAB units, the RGB and Opponent Colours based systems rated considerably better than the other selection systems during the first session half. Although this superiority in performance persisted in the second session half, the differences between the RGB and Opponent Colours based systems and the other systems diminished significantly.

Thus, there is some evidence here to suggest that subjects who used the RGB or Opponent Colour based systems performed well (according to the criteria mentioned above) at the *beginning* of the session as well as later on in the experiment. On the other hand, subjects using systems based on HSV, YIQ or CIELAB seemed to encounter some difficulty at the start of the sessions. After a moderate amount of practice the performance of these subjects improved considerably but was still not as good as the performance of the subjects who used the RGB or Opponent Colours systems.

Whether the apparent lack of learning in the performance of the subjects who used the RGB or Opponent Colours systems is an advantage or disadvantage is not entirely clear. For instance, it may be that these systems are very "natural" for inexperienced users to select colours with. On the other hand, it is possible that it is simply very difficult for subjects to enhance their performance using these systems and that the other systems may become easier to use for colour selection after extensive usage. Thus it is important to take into account here, as well as for the other conclusions about learning that are derived from the performance of the inexperienced subjects, that these subjects "practiced" for only about an hour (the approximate length of a typical colour matching session). There are indications, however, that the RGB and Opponent Colours selection systems are, in fact, more "natural" to use than the others. A thorough discussion of this hypothesis is deferred until later in this chapter.

7.1.5. No Significant Differences Between Single and Multiple Attribute Control

In the data presented in the previous chapter, there do not appear to be any obvious behavioral trends that would indicate any significant differences between the performance of subjects using colour selection systems that allowed simultaneous control of multiple colour attributes and those subjects that used systems where it was only possible to control one colour attribute at a time. Where any performance differences did occur they tended to be specific to one or two performance measures. Furthermore, there were no instances of significant differences between the two interaction techniques that persisted across all colour order systems for any one particular performance measure.

Hence, there does *not* appear to be any strong evidence in the results from this experiment indicating that subjects who used selection systems based on "multiple attribute control" performed significantly different from subjects who performed the colour matching tasks using systems based on "single attribute control."

Based on this observation, it is apparent then that allowing users to control more than one colour attribute simultaneously does not, in general, significantly improve or degrade their ability to perform colour selection tasks. Some possible explanations for this observation are discussed in Section 7.3.2.

7.1.6. Environments With Very Little Ambient Light Have Undesirable Side Effects

In Section 6.3, a summary of the comments made by the inexperienced subjects at the completion of each session was given. From Table 6.3 it can be seen that several subjects commented that the level of light emitted from the display devices was too high or that the level of light in the experiment room was too low. In addition, a number of subjects complained that they suffered from eye strain during the experiment or that the experiment lasted too long.

It is apparent from these comments that the very low level of ambient light in the experiment room produced a sharp contrast between the light emitted from the CRT display devices and the relative darkness of the surrounding environment. This in turn appeared to cause minor visual fatigue problems when subjects had to perform visually strenuous tasks for moderately long periods of time (i.e. approximately 1 hour).

In general, then, situations where users of computer equipment are required to work for long periods of time in an environment of very low level of ambient light should be avoided. This is especially applicable in "professional" environments where the extensive use of CRT-type display devices is very common.

7.2. Models of Subject Performance

A behavioral model is developed in this section in order to partition the process of colour matching according to the performance patterns observed in the data collected during the experiment. A number of conclusions are then derived from the results using this model.

7.2.1. Colour Matching in Terms of Phases of Subject Performance

When analyzing the results of a psychophysical experiment such as this one, it is often useful to be able to subdivide the observed interactive process (in this case the process of colour matching) into parts, or *phases*. One method of achieving such a partitioning is to examine the results, both at the level of individual trials and in the aggregate across all trials, and determine if there are any significant changes in performance as subjects get closer to accomplishing the required task (i.e. achieving a colour match). For example, do subjects seem to perform differently near the completion of the colour matching task than they do at the beginning of a trial? If such transitions in observed behavior are apparent, then the process can be subdivided into phases according to where or when in the execution of the task the performance of subjects appears to change. Subsequent analysis of the data can then be carried out on the basis of this partitioning, using the different (and similar) performance characteristics of the various task subdivisions to further develop conclusions about how subjects went about accomplishing the required task.

It was initially hypothesized that the colour matching process could be subdivided into (at least) two distinct phases since, subjectively, one could sense that the strategy used to perform the colour matching task seemed to change as one got closer to the target colour. This conjecture was reinforced when informal inspection of the records of individual colour matching trials revealed a pattern in the data that seemed to recur with significant regularity.

This pattern can best be described using the data of a "typical" colour matching trial. Figure 7.1 traces the difference in colour between the two test stimulus subfields for the duration of such a typical trial. Colour difference is plotted in the vertical direction using CIELAB colour difference units, while the horizontal axis represents time (in seconds) from the start of the trial.

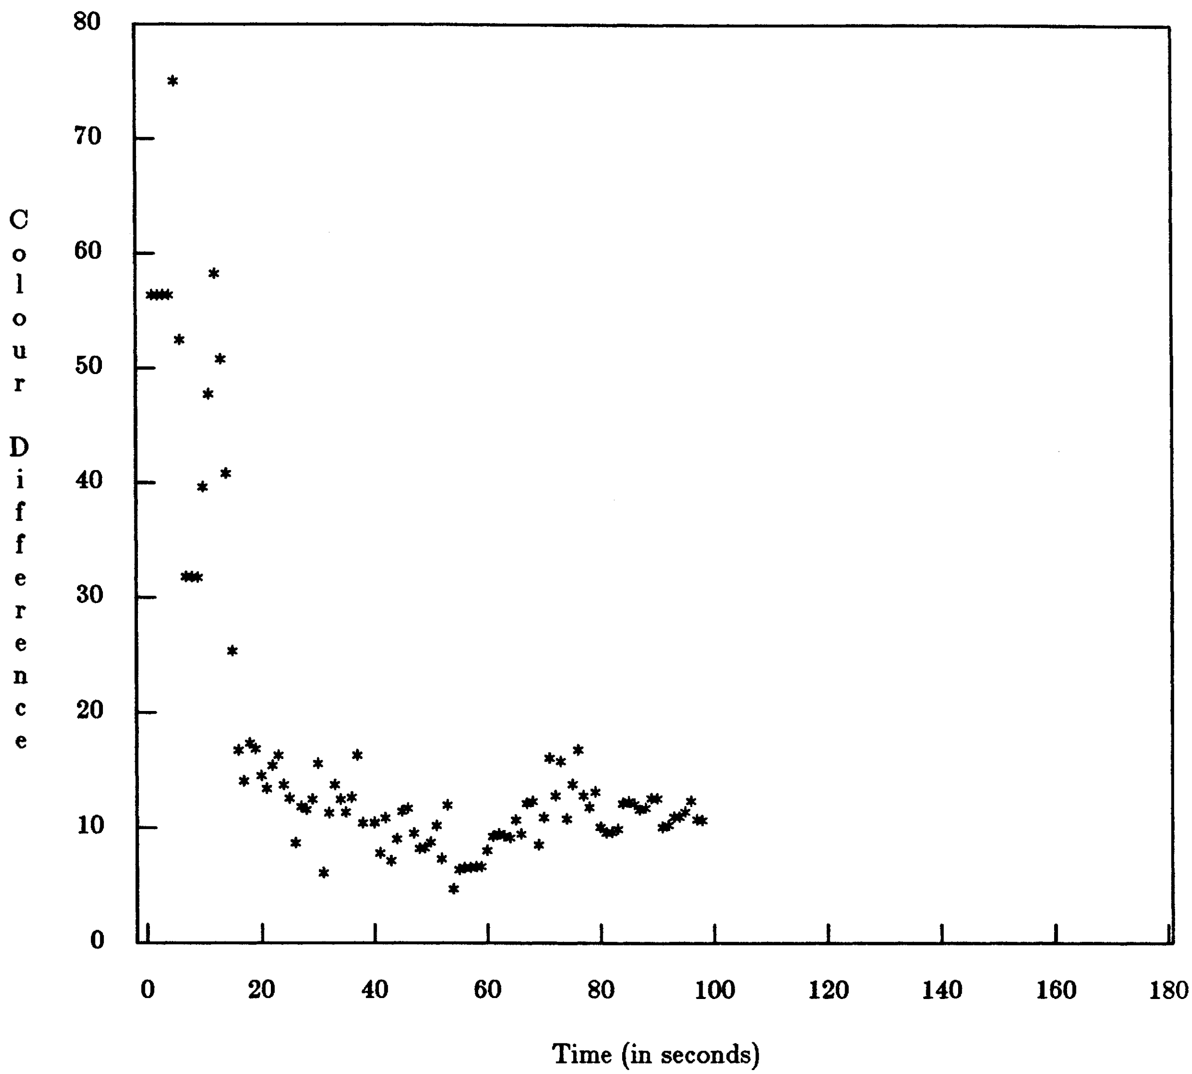

Figure 7.1. Path of colour difference for the duration of a "typical" colour matching trial.

Two distinct phases are apparent in this data: an early phase (subsequently referred to as the *convergence phase*) that is characterized by a rapid decrease in the colour difference measure and a later phase (subsequently referred to as the *refinement phase*) that is marked by small, almost random, fluctuations in colour difference that occur at a level relatively close to the X axis (i.e. a "perfect" match). Intuitively, this behavioral pattern seems to indicate that, in a "typical" trial, the subject was at first able to manipulate the colour of the controllable subfield in the "direction" of the target colour quite rapidly. Once the controllable colour was close to the target, the subject seemed to reach a stage where it was much more difficult to converge the colours at such a rapid rate and a considerable amount of time was spent making rather detailed refinements to the match.

It should be pointed out that, even though the colour difference measure remains at a more or less constant level in the later stage (i.e. the refinement phase), this does not necessarily imply that the colour controlled by the subject was not being modified. The difference measure plotted in the figure is an *absolute* measure of colour difference, meaning that many different colours of the controllable subfield will result in an identical colour difference value (i.e. all these colours are the same distance from the target). Thus in many instances where the difference values remain relatively constant across time, it may very well be the case that the colour of the controllable subfield was substantially modified, but remained, nevertheless, at almost the same distance from the target colour (e.g. the controllable colour may have been "circling" around the target). Also, since the test stimulus was sampled at discrete intervals, it is possible that the controllable colour was brought substantially closer to the target (or even right to the target) between sample points.

The data shown in Figure 7.1 is presented here as being representative of hundreds of colour matching trials. Although visual examination of the records of many trials showed that the pattern illustrated in the figure was followed to some extent in most colour matches, it was not entirely clear how the effects of this pattern could be measured in the aggregate. Ideally, the precise point of transition between the two phases could have been determined for each trial. A number of useful performance measures could then have been calculated based on this data (e.g. the mean time taken to reach the refinement phase and the mean colour difference at the end of the convergence phase). The main problem with this strategy was that, although the colour matching records frequently showed a definite transition between phases, there did not appear to be any consistent point in a trial at which this transition took place. Nor was there any pattern in the level of colour difference at which it occured. Furthermore, the visual appearance of the phases varied considerably from trial to trial (e.g. in slope and smoothness). As a consequence, it became evident that to accurately determine where the transition point between phases occured in the general case would require the development of extremely sophisticated pattern matching algorithms.

Since performing such a complex analysis was generally beyond the scope of the investigation, it was apparent that a less precise method would have to suffice. The technique that seemed most appropriate was based on using colour difference thresholds to examine the colour matching process at different stages of task completion. As was described in Section 6.1.5, a colour difference threshold is a level of colour matching precision defined in CIELAB colour difference units. A threshold is "reached" at the point in a trial where colour difference is measured to be less than or equal to the threshold value for the first time in the trial.

By calculating the average time it took subjects to reach various threshold levels, the performance of subjects during the different phases of colour matching could be examined. Since relatively high levels of colour difference were typically reached relatively early in the matching process (i.e. during the convergence phase or in the early part of the refinement phase according to the pattern illustrated in Figure 7.1), a large proportion of the *aggregate* effort used to reach these thresholds will have been spent in the convergence phase. Thus the average (calculated across a particular subset of trials) time taken to reach relatively high colour difference thresholds will measure, for the most part, the performance of subjects during the period of relative rapid convergence (i.e. the convergence phase). At lower levels of colour difference, a larger proportion of the effort exerted to reach the thresholds will occur after the period of rapid convergence (i.e. during the refinement phase). Hence, the measurement of the average time taken to reach these thresholds will contain a large component of activity that occured during the refinement phase. The lower the level of colour difference, the larger this component will be and consequently the more the performance measure will indicate performance during the refinement phase.

In order to determine which thresholds of colour difference sufficiently characterize the performance of inexperienced subjects during the two phases, performance at a total of 15 threshold levels was examined. As was noted in Section 6.1.5, a definite trend could be observed in the data from these thresholds. At relatively high levels of colour difference, subjects who performed the experiment using colour selection systems based on the RGB or Opponent Colours models were able to reach the thresholds significantly faster than subjects who used other selection systems. At low colour difference thresholds, this apparent advantage of the RGB and Opponent Colours based systems disappeared; that is, the subjects who used these systems did not perform better than other subjects.

A subset of 3 of the 15 threshold levels (i.e. 25, 14 and 6 CIELAB colour difference units) were chosen to represent the extremes and median of this transition in subject performance. The average times taken by subjects to reach these levels of colour difference using different selection systems was presented in Section 6.1.5. The superior performance of the RGB and Opponent Colours systems is quite clear at a threshold of 25 CIELAB units. At 14 CIELAB units, subjects who use these selection systems are still able to reach the threshold faster than subjects who used other systems but the magnitude of the performance differences is considerably less than at the 25 CIELAB level. At a level of 6 CIELAB units, the advantage of the RGB and Opponent Colours systems has vanished altogether.

Since it is evident that the performance of subjects changes significantly between the higher threshold levels and the lower ones, this suggests that there may be a direct link between the distinct phases observed in many of the records of individual trials and the observation that subjects were able to reach thresholds of 25 and 14 CIELAB units faster when using RGB and Opponent Colours based selection systems but could not do so at lower threshold levels. Although this connection is difficult to verify, it does seem intuitive that there is in fact a correlation between the two observed phenomena. The model presented here is therefore based on the premise that the performance advantage exhibited by the RGB and Opponent Colours based selection systems at higher threshold levels is a property of performance during the convergence phase. Conversely, approximatelly equal performance at lower threshold levels is a property of subject performance during the refinement phase.

According to the model then, the colour difference threshold of 25 CIELAB units is used to represent the behavior of subjects performing colour matching during the convergence phase. At the 6 CIELAB unit level of colour difference, the observed phenomena are interpreted as having occurred at a stage in the colour matching process that is largely dominated by the performance behavior exhibited during the refinement phase. The 14 CIELAB colour difference threshold shows a mixture of performance characteristics from both phases.

The remainder of this section discusses a number of conclusions based on this model.

7.2.2. RGB and Opponent Colours Show Superior Performance During Convergence Phase

According to the model developed above, inexperienced subjects were able to perform colour matching faster during the convergence phase using selection systems based on the RGB or Opponent Colours models than they could using other selection systems. This statement is reinforced by the observation that the experienced subjects also seemed to perform better during the convergence phase when using these systems (as is indicated by the data for colour difference thresholds of 25 and 14 CIELAB units presented in Section 6.4.4).

From these observations it is possible to predict that, for colour selection tasks where it is only important to select a colour in the general vicinity of a "target" colour (i.e. detailed refinement of the colour is not important), inexperienced users (and possibly experienced users as well) will perform the tasks faster with colour selection systems based on the RGB or Opponent Colours models than with selection systems based on the HSV, CIELAB or YIQ colour models.

7.2.3. No Clear Optimal Systems For Refinement Phase

For the refinement phase, there did not appear to be any clear patterns in the results that would indicate which of the colour selection systems would be best for inexperienced users. For example, in the data for the colour difference threshold of 6 CIELAB units where a lot of the matching activity before reaching the threshold was spent in the refinement phase, none of the selection systems stood out in terms of this criteria. Any performance differences that did occur tended to be much less significant than in the convergence phase.

Hence, it appears that the colour order system with which colour selection tasks are performed does not have a significant impact on performance during the refinement phase. As a consequence, there does not appear to be a colour order system that is clearly "optimal" in allowing users to easily perform colour selection tasks where it is important to have "fine" control over the colours used in an image.

7.2.4. Learning May Affect Convergence Phase More Than Refinement Phase

In Section 6.2, a number of performance measures were averaged by "colour slice" across all colour matching trials. One way to examine how the performance of subjects changed as they became familiar with the selection system they were using is to compare average performance during the first colour slice to performance during the last colour slice. As can be seen in Figures 6.7a and 6.8a, however, the performance measures shown there seem to take unusual jumps between the fifth and sixth colour slices. One possible reason for these jumps may be that subjects were becoming fatigued at this stage in the experiment and were anxious to finish off the last set of trials. In any case, the fifth colour slice is used here for comparative purposes rather than the last one in order to avoid these anomalies. Furthermore, all the calculations quoted below are based on averages that include incomplete trials but do not include trials where the white target colour was matched (since these trials start at a colour difference that is less than 25 CIELAB units).

Of particular interest here is the effect that "learning" has in the different phases of task performance. The mean time taken by subjects to reach a colour difference threshold of 25, 14 and 6 CIELAB units decreases by 31%, 39% and 33% respectively between the first and fifth colour slices. On the other hand, the mean total time taken by subjects to complete the colour matching task decreases by only 18% between the first and fifth colour slices.

One immediate inference is that for *each* of the colour difference thresholds included in the current analysis, the proportion of time spent by subjects in reaching a particular threshold (of the total time to complete the colour matching task) was considerably less in the fifth colour slice than in the first colour slice. Hence, the proportion of the total matching time needed to reach a particular colour difference threshold seems to decrease with practice.

The practical implications of this, however, are not entirely clear. The significant differences between the large decreases in time spent reaching the 25 and 14 CIELAB unit thresholds and the relatively small decreases in the total time spent matching would seem to lead to the conclusion that most of the improvement in performance takes place in the convergence phase. However, if this were indeed true, then one would also expect that the decrease in mean time to reach a threshold would be significantly less at the 6 CIELAB unit level than at the 25 and 14 CIELAB unit level. Since this is not the case, it may be that there are other phenomena occuring here that are not evident from the current analysis.

Nevertheless, there is some evidence here that suggests that most learning does in fact occur in the convergence phase. Given this tentative conclusion, it would seem that the implication of this for colour selection is as follows. As users become more familiar with a colour selection system, their performance in selecting a colour "in the general vicinity" of a desired target colour will improve at a faster rate than their performance in selecting a colour that is very close to a particular desired colour.

7.2.5. Two Types of Experience Have Equivalent Advantages During Convergence Phase

As was noted in Chapter 6, the two experienced subjects who took part in the experiment had rather different forms of experience. In one case, the subject had spent many years performing colour matching tasks using the RGB colour model but had never used a tablet or mouse before. The other subject had a lot of experience in using tablets but had relatively little experience in performing colour selection tasks.

Not surprisingly, the two experienced users performed moderately better than the averages for the inexperienced users according to many criteria. For example, both of these subjects were, on average, as fast or faster than the inexperienced subjects in reaching colour difference thresholds of 25, 14 and 6 CIELAB units for every one of the selection systems that were used by the experienced subjects.

One observation that is a little unusual is that the subject that had significant experience using tablets was able to reach the 25 and 14 CIELAB unit thresholds using RGB based selection systems just as quickly as the subject having significant experience matching colours using the RGB colour model. At the 6 CIELAB unit level, however, the subject with experience using RGB did seem to have an advantage in performance.

Thus, there is some evidence here that suggests that, at least in the convergence phase, the advantages in performance gained through significant familiarity with a particular colour order system appear to be approximately equivalent to the advantages gained through significant experience in using a particular interaction device.

7.2.6. Familiarity With Colour Model Used May Also Help in Refinement Phase

In the results discussed so far, there has not been any evidence to suggest that the colour order system used to perform colour selection will give the user any advantage during the refinement phase. From the results of the experienced subjects presented in Section 6.4, however, a number of interesting patterns can be seen that indicate a rather unique behavior for Subject 1 during the refinement phase. The reader will recall that this subject has participated in colour matching experiments similar to this one for approximately 30 years. In fact, all the colour matching operations performed in these experiments were done using selection schemes based on the RGB colour model. Thus, this subject had extensive experience with RGB based colour selection systems but had no experience using any of the other colour order systems evaluated here.

From Figure 6.16a it can be seen that Subject 1 reached a threshold of 6 CIELAB units significantly faster, on average, using the RGB based selection systems than with systems based on the other colour models. The differences between systems in the mean time taken to reach this threshold were not as significant as at the 25 and 14 CIELAB unit levels. However, the proportion of trials that actually reached this closeness threshold was at least 80% for both of the RGB systems. This is considerably more than for the other selection systems. Furthermore, the mean final colour difference measure (as shown in Figure 6.13a) was lower for the RGB based systems than for all but one of the systems based on other colour models.

Thus, there is some evidence to suggest that having extensive experience in performing colour matching operations using a particular colour order system seems to moderately enhance performance in the refinement phase when the task is performed using the familiar colour model. Hence, a user will be able to select colours that are "very close" to a desired target colour moderately faster using a selection system that is based on a colour model with which the user is very familiar (i.e. a model that has been extensively used by the user for performing colour selection).

7.3. Possible Explanations for Observed Behavior

In this section, a discussion of possible explanations for the performance behavior observed in the colour matching experiment is given. The arguments presented are speculative since no effort was made to support them with empirical evidence. Nevertheless, because of the interesting implications of these hypotheses, the author believes that it would be worthwhile to conduct further investigations to determine to what extent the proposed explanations account for the obeserved phenomena.

7.3.1. Why RGB and Opponent Colours Systems Appear To Be "Natural" To Use

The discussion in the previous two sections has pointed out that, as a group, the colour selection systems based on the RGB and Opponent Colours colour models seem to stand out in terms of task performance during the early stages of the colour matching process (i.e. the convergence phase). Furthermore, it appears that a moderate amount of practice does not significantly affect this performance. These observations suggest that these colour models may, in fact, have inherent properties that are quite "natural" for users to relate to when performing interactive colour selection tasks. If this is indeed true, then it is quite possible that these models have properties that are close to the characteristics of one of the "hardwired" psychological processes of the human visual system.

As was explained in the description of the Opponent Colours colour order system in Chapter 4, one such process is believed to exist at an intermediate level in the hierarchy of the visual system. The Opponent Colours theory states that after light is initially absorbed by the cones in the retina, neural signals from the three cone types are transmitted via separate channels to a second processing stage. The output of this process is three new signals: an achromatic (or luminance) signal, a signal that differences red-green colour information, and a signal that differences yellow-blue colour information. Thus, of the chromatic properties of visible light, the red, green, blue and yellow attributes are believed to be fundamental to the way humans perceive colour. The selection systems based on both the RGB and Opponent Colours models have independent degrees of freedom that allow users to control attributes that are among these fundamental chromatic properties. Whether this is in fact the real reason why subjects using these selection systems exhibited superior performance is mostly a matter of speculation at this point but the explanation does have intuitive appeal.

7.3.2. Why the Number of Attributes Controlled Simultaneously Doesn't Matter

According to the results reported, there did not seem to be any significant differences between the performance of subjects who were permitted to manipulate two degrees of freedom simultaneously and those subjects who were required to alter each degree of freedom separately. An explanation for this phenomenom may be found in the strategy subjects employed in using the "multiple attribute control" mechanism. For instance, one of the experienced users remarked that when he was using a selection system that allowed the simultaneous manipulation of two colour properties, he tended to move the puck in an exactly vertical or horizontal direction so that only one of the properties were actually being altered at any one time. Hence, the facility for updating two degrees of freedom at once was generally ignored and the subject consequently tended to emulate the "single attribute control" mechanism.

One possible reason for such behavior may be that the instructions given to subjects prior to performing the experiment indirectly influenced how subjects used the multiple attribute control mechanism. Both the "box" analogy used in the instructions to emphasize the three dimensional nature of colour and the technique of introducing the controllable colour attributes one at a time may have implicitly encouraged subjects to alter the control properties in a sequential manner. On the other hand, some effort was made in the wording of the instructions to encourage subjects to exploit the ability to alter two degrees of freedom simultaneously. It may have been the case that subjects simply found it easier to become familiar with the colour selection systems by changing one colour attribute at a time. Further insight in this regard might be gained by comparing the performance of subjects who used multiple attribute control selection systems in this experiment to the performance of subjects who are explicitly encouraged to make colour manipulations by simultaneously altering two degrees of freedom (e.g. using "circular" puck motions to converge upon the target colour).

If the behavior described by the experienced subject is characteristic of the performance of other subjects, then it would not be surprising if subjects using the two different interaction techniques did not perform all that differently. However, from the summary of comments presented in Section 6.3, it can be seen that there was in fact only one subject from the inexperienced category that mentioned that this strategy had been used. On the other hand, this type of comment is not necessarily something that one would expect subjects to frequently comment about even if this method of interaction was frequently used. Again, in lieu of a more thorough examination, the explanation presented here should be treated as speculation rather than as a statistically supported hypothesis.

8. Future Research

Although a number of behavioral patterns have become apparent from the results of the current investigation, it is clear that the present understanding of how users of interactive computer graphics systems perform colour selection tasks is still not complete. This study has provided several insights into the problems involved but many questions still remain unanswered. In addition, a number of new issues have been raised. In this chapter, some direction for future investigations in this area of research is provided.

8.1. Follow-up Investigations

The results of the colour matching experiment suggest a number of useful follow-up studies:

- One of the key patterns observed in the empirical data was that, during the initial stages of colour matching, subjects who used selection systems based on the RGB or Opponent Colours model seemed to have performed considerably better than subjects who used other systems. Most of the target colours that subjects were required to match, however, tended to be relatively closely aligned with the colour properties that subjects were able to directly control when these colour models were used. For instance, one could argue that subjects who used RGB could easily get close to the red target colour by manipulating only one degree of freedom. In general, the required task was not as simple as this since a deliberate effort was made to choose target colours that made the task more difficult. Performing a similar experiment using other target colours that seem intuitively more difficult to match (e.g. yellow met this criterion in the case of the RGB model) might provide additional insight into the issues involved in performing these tasks. In particular, there were indications in the data that suggested that subjects who used the RGB systems encountered more difficulty in matching white than in matching the other target colours. Hence, including some colours of intermediate saturation might prove useful.
- The Opponent Colours model implemented for evaluation in the current experiment was designed as an initial attempt at emulating the fundamental properties of the theory upon which this system was based. Since the behavior of subjects using the Opponent Colours selection systems seemed to indicate several desirable characteristics, it would appear that evaluating more advanced versions of this colour model would be quite useful in providing a further understanding of the phenomena observed in this experiment. For example, modifications to the implemented system could be made that reflect some of the enhancements to the Opponent Colours model that have been suggested in the colour vision literature. In addition, an implementation that conforms more regularly to the colour gamut of CRT

display devices would be desirable. This would allow subjects to more easily be able to predict the locations of the "boundaries" of the displayable colour space .

- A more comprehensive examination of the learning component would shed additional light on which systems would be most effective for extensive usage. In the present experiment, the period of observable practice time was limited to approximately one hour. Conducting further studies where subjects use a fixed colour selection system for considerably longer periods of time might resolve, for example, the issue of whether extensive practice has an effect on performance during the refinement phase.
- Examining the performance of different "classes" of experienced users might be useful for providing hints on how to design colour selection interfaces for specific applications. For instance, a study of artists performing colour matching might be of interest to designers of interactive picture creation systems.
- There was some suggestion in the current experiment that a subject's perception of relative task performance did not always conform with what was indicated by the objective measures of performance. In other words, the selection system a subject felt was "best" to use was not necessarily the most optimal according to the performance measures employed. Thus, it would be of great interest to measure the extent to which this phenomena occurs in general and also to determine which components of the colour matching process contribute to a subject's perception of task difficulty and which do not. One method of doing this might be to devise an experiment that more closely resembles "real world" colour selection tasks so that decisions on whether a selected colour is "good enough" are driven more by a subject's judgement of aesthetics rather than by comparison to a preselected target colour. These results could then be compared to the results from the colour matching experiment to gain insight on the differences between subjective and objective performance. On the other hand, designing meaningful subjectively-driven experiments is, in general, extremely hard. Some of the difficulties in constructing such experiments so that performance can be suitably measured were discussed in the body of this thesis.

8.2. Further Data Analysis

Although a major effort was made in analyzing the large volume of data produced from the current experiment, several issues of task performance were omitted from the analysis. A number of suggestions for further analysis of this data are presented below.

- Determine to what extent subjects using selection systems based on the "multiple attribute control" mechanism performed the colour matching tasks in a manner that emulated "single attribute control." This can be achieved in a relatively straightforward manner by examining the colour matching records for colour manipulations that occurred during the times in which the puck buttons that allowed simultaneous control of two degrees of freedom were depressed.

- Investigate in more detail the actions of subjects during the refinement phase of colour matching. Is it valid, for instance, to consider the colour manipulations that occurred during these later stages as being *random* as far as task performance is concerned? Is there any correlation at all between the time a subject spends "fine tuning" a colour and the actual closeness of the match?
- Develop additional behavioral models in order to better comprehend the interactions observed. One such model might, for example, express task performance in terms of a bounding volume that surrounds the target colour in colour space. Within the volume, modifications to the colour of the test stimulus might be considered random, while for colour manipulations outside this threshold the controllable colour is deterministically altered so that the colour quickly moves inside the bounding volume again. The performance of subjects using a particular selection system could then be inversely related to the magnitude of the bounding volume.
- Investigate whether the number of times subjects switched between degrees of freedom by pressing different puck buttons had any noticeable effect on task performance. This measure would indicate what the relationship between "mode switching" and subject performance is.
- Perform a more thorough analysis of which colour models were most effective (or least effective) in matching a given target colour. For example, with which selection systems were subjects able to match white more easily. Some preliminary analysis in this regard was performed in the current investigation, but it was evident from the results obtained (which were not included in this report) that more work would have to be done before any conclusive statements can be made.

8.3. Other Issues Affecting Colour Selection

The current experiment was designed primarily to evaluate a number of colour order systems and interaction techniques in terms of how suitable each is for performing interactive colour selection tasks. These components of the human-computer interface are just two of many factors that can affect user performance however. Some suggestions to shed some light on the impact that other factors have on user performance are given below.

- Examine the effect, if any, that the use of different control devices has on the performance of colour selection tasks. In particular, an analysis of differences in performance between position sensitive devices (e.g. sliders, knobs, tablets used in the normal position sensitive manner) and motion sensitive devices (e.g. mice, trackballs) might produce some useful results.
- Analyze the impact that various forms of feedback have on the performance of colour selection tasks. For example, visual or aural cues to indicate that the boundary of the displayable colour gamut has been encountered, feedback mechanisms that show what the current level of each controllable colour attribute is, or iconic visual aids that help the naive user

learn which physical degrees of freedom control which perceptual colour properties might prove to be useful in making it easier for a user to select colours. In addition, it would be valuable to determine exactly what colour updating rates are necessary to prevent the user from perceiving discrete jumps in colour.

- Investigate whether combinations of different colour selection techniques can make colour selection tasks easier to perform. One such interface that has some intuitive appeal allows the user to initially select a colour from a palette. To make finer adjustments to this colour, the user could then invoke a selection system similar to those evaluated in the present experiment.

Appendix A

Implementation Details of Colour Order Systems

This appendix contains the implementation details of the colour order systems evaluated in the colour matching experiment. The implementation of these colour models is based on a hierarchy of coordinate transformations as shown in Figure A.1. These transformations are used to successively convert colour coordinates specified in systems found higher up in the hierarchy into coordinates of the lower systems. Each colour specified in any of the colour order systems is ultimately transformed into the device dependent coordinates used by the computer graphics hardware. In the case of the graphics system used to run this experiment (i.e. the Adage / Ikonas RDS 3000 system), the low level colour information is represented as three 10-bit entries in a colour lookup table. These 10-bit values are translated by digital-to-analog converters into the driving voltages of the red, green and blue electron guns of the CRT display device. Each of the colour models implemented were transformed into these device level coordinates directly from the RGB colour system. The "gamma correction" algorithm used to achieve the transformation from the RGB system to device coordinates is well documented in [Cowan83a] and is therefore not repeated here. The other transformations indicated in Figure A.1 are described in the remainder of the appendix.

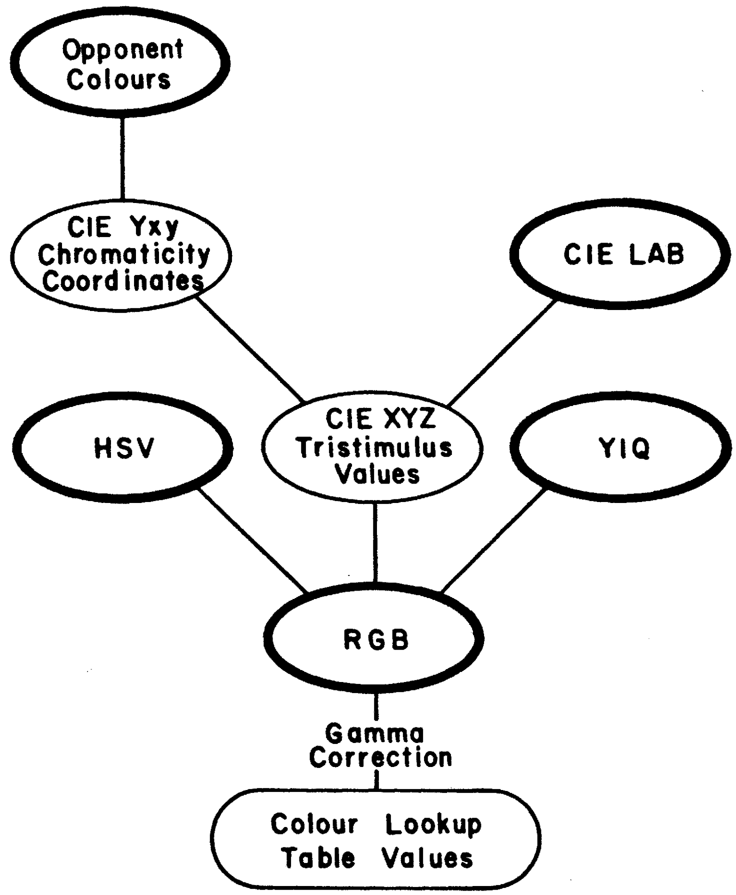

Figure A.1. Hierarchy of colour transformations used to reduce colours specified in coordinates of the colour order systems to device coordinates.

HSV to RGB Transformations

(R,G,B) <- (H,S,V):

```
F  :=  H * 6 - floor (H * 6);
T1  :=  V * (1.0 - S);
T2  :=  V * (1.0 - (S * F));
T3  :=  V * (1.0 - (S * (1.0 - F)));

case   (floor (H * 6)) mod 6   of
        0:  (R,G,B)  :=  (V,T3,T1);
        1:  (R,G,B)  :=  (T2,V,T1);
        2:  (R,G,B)  :=  (T1,V,T3);
        3:  (R,G,B)  :=  (T1,T2,V);
        4:  (R,G,B)  :=  (T3,T1,V);
        5:  (R,G,B)  :=  (V,T1,T2);
end
```

(H,S,V) <- (R,G,B):

```
MIN  :=  min (R,G,B);
MAX  :=  max (R,G,B);
SPREAD  :=  MAX - MIN;

V  :=  MAX;
if  V = 0.0  then
begin
        H  :=  0.0;
        S  :=  0.0;
        return;
end

S  :=  SPREAD / MAX;
if  S = 0.0  then
begin
        H  :=  0.0;
        return;
end

R'  :=  (MAX - R) / SPREAD;
G'  :=  (MAX - G) / SPREAD;
B'  :=  (MAX - B) / SPREAD;
if  MAX = R  then
        H  :=  (if  MIN = G  then  5.0 + B'  else  1.0 - G');
else  if  MAX = G  then
        H  :=  (if  MIN = B  then  1.0 + R'  else  3.0 - B');
else
        H  :=  (if  MIN = R  then  3.0 + G'  else  5.0 - R');
H  :=  H / 6.0;
```

<u>YIQ to RGB Transformations</u>

(R,G,B) <- (Y,I,Q):

$$\begin{pmatrix} R \\ G \\ B \end{pmatrix} := \begin{pmatrix} 1.00 & 0.95 & 0.62 \\ 1.00 & -0.28 & -0.64 \\ 1.00 & -1.11 & 1.73 \end{pmatrix} \cdot \begin{pmatrix} Y \\ I \\ Q \end{pmatrix}$$

(Y,I,Q) <- (R,G,B):

$$\begin{pmatrix} Y \\ I \\ Q \end{pmatrix} := \begin{pmatrix} 0.30 & 0.60 & 0.21 \\ 0.59 & -0.28 & -0.52 \\ 0.11 & -0.32 & 0.31 \end{pmatrix} \cdot \begin{pmatrix} R \\ G \\ B \end{pmatrix}$$

<u>CIEXYZ to RGB Transformations</u>

N.B. In the following transformation pair, the symbols x_R, y_R, z_R, x_G, etc. represent the CIEYxy chromaticity coordinates of the red, green and blue phosphors of the CRT display monitor used to conduct the experiment. The values of these chromaticities are given in Table A.1. of this appendix.

(R,G,B) <- (X,Y,Z):

$$\begin{pmatrix} R \\ G \\ B \end{pmatrix} := \begin{pmatrix} x_R & x_G & x_B \\ y_R & y_G & y_B \\ z_R & z_G & z_B \end{pmatrix}^{-1} \cdot \begin{pmatrix} X \\ Y \\ Z \end{pmatrix}$$

(X,Y,Z) <- (R,G,B):

$$\begin{pmatrix} X \\ Y \\ Z \end{pmatrix} := \begin{pmatrix} x_R & x_G & x_B \\ y_R & y_G & y_B \\ z_R & z_G & z_B \end{pmatrix} \cdot \begin{pmatrix} R \\ G \\ B \end{pmatrix}$$

CIELAB to CIEXYZ Transformations

(X,Y,Z) <- (L,A,B):

$$Y^* := \frac{L + 16}{116}$$

$$X := \left(\frac{A}{500} + Y^*\right)^3$$

$$Y := \left(Y^*\right)^3$$

$$Z := \left(-\frac{B}{200} + Y^*\right)^3$$

(L,A,B) <- (X,Y,Z):

$$L := 116 * Y^{1/3} - 16$$

$$A := 500 * (X^{1/3} - Y^{1/3})$$

$$B := 200 * (Y^{1/3} - Z^{1/3})$$

CIEYxy to CIEXYZ Transformations

N.B. In order to distinguish the Y coordinate of the CIEYxy system from the Y tristimulus value, the symbol L is used here to represent the CIEYxy Y value.

(X,Y,Z) <- (L,x,y):

X := L * x / y;
Y := L;
Z := L * (1.0 - (x + y)) / y;

(L,x,y) <- (X,Y,Z):

L := Y;
x := X / (X + Y + Z);
y := Y / (X + Y + Z);

Opponent Colours to CIEYxy Transformations

N.B. Several constants are used in the transformations given below. These constants are defined as follows:

WHITE_x, WHITE_y	CIEYxy chromaticity coordinates of the white reference point.
THETA_RED	Angular coordinate of the red reference point with respect to the white reference point in CIEYxy space.
THETA_YELLOW	Angular coordinate of the yellow reference point with respect to the white reference point in CIEYxy space.
THETA_GREEN	Angular coordinate of the green reference point with respect to the white reference point in CIEYxy space.
THETA_BLUE	Angular coordinate of the blue reference point with respect to the white reference point in CIEYxy space.
THETA_RY	Angular difference between THETA_YELLOW and THETA_RED.
THETA_YG	Angular difference between THETA_GREEN and THETA_YELLOW.
THETA_GB	Angular difference between THETA_BLUE and THETA_GREEN.
THETA_BR	Angular difference between THETA_GREEN and (THETA_RED + 2*PI).
HALF_PI	Pi divided by 2.

The CIEYxy chromaticity coordinates of the white, red, green, blue and yellow reference points are given below in Table A.1.

Reference Colour	Chromaticity Coordinates		
	x	y	z
White	.31	.32	.37
Red	.62	.33	.05
Green	.21	.68	.11
Blue	.15	.06	.79
Yellow	.51	.49	.00

Table A.1. Chromaticity coordinates of the reference colours used in the implementation of the Opponent Colours model. The red, green and blue reference colours are of the same chromaticity as the phosphors of the CRT display monitor used in the experiment.

(Y,x,y) <- (LUM,RG,YB):

```
RADIUS := sqrt (RG*RG + YB*YB);

if YB >= 0.0 and RG >= 0.0 then
begin
      /* Yellow - Green Quadrant */
      THETA_OPP := arctan (RG / YB);
      THETA_Yxy := THETA_YELLOW + THETA_OPP / HALF_PI * THETA_YG;
end

if YB < 0.0 and RG >= 0.0 then
begin
      /* Green - Blue Quadrant */
      THETA_OPP := arctan (-YB / RG);
      THETA_Yxy := THETA_GREEN + THETA_OPP / HALF_PI * THETA_GB;
end

if YB < 0.0 and RG < 0.0 then
begin
      /* Blue - Red Quadrant */
      THETA_OPP := arctan (RG / YB);
      THETA_Yxy := THETA_BLUE + THETA_OPP / HALF_PI * THETA_BR;
end

if YB >= 0.0 and RG < 0.0 then
begin
      /* Red - Yellow Quadrant */
      THETA_OPP := arctan (YB / -RG);
      THETA_Yxy := THETA_RED + THETA_OPP / HALF_PI * THETA_RY;
end

Y := LUM;
x := WHITE_x + RADIUS * cos (THETA_Yxy);
y := WHITE_y + RADIUS * sin (THETA_Yxy);
```

(LUM,RG,YB) <- (Y,x,y):

```
/* Convert relative (x,y) values to polar coordinates. */
RELX := x - WHITE_x;
RELY := y - WHITE_y;
RADIUS := sqrt (RELX*RELX + RELY*RELY);
THETA_Yxy := arctan (RELY / RELX);

/* Yellow - Green Quadrant */
if THETA_Yxy >= THETA_YELLOW and THETA_Yxy < THETA_GREEN then
    THETA_OPP := (THETA_Yxy - THETA_YELLOW) / THETA_YG * HALF_PI;

/* Green - Blue Quadrant */
if THETA_Yxy >= THETA_GREEN and THETA_Yxy < THETA_BLUE then
    THETA_OPP := (THETA_Yxy - THETA_GREEN) / THETA_GB * HALF_PI
                 + HALF_PI;

/* Blue - Red Quadrant */
if THETA_Yxy >= THETA_BLUE then
    THETA_OPP := (THETA_Yxy - THETA_BLUE) / THETA_BR * HALF_PI
                 + PI;
if THETA_Yxy < THETA_RED then
    THETA_OPP := (THETA_Yxy + 2*PI - THETA_BLUE) / THETA_BR * HALF_PI
                 + PI;

/* Red - Yellow Quadrant */
if THETA_Yxy >= THETA_RED and THETA_Yxy < THETA_YELLOW then
    THETA_OPP := (THETA_Yxy - THETA_RED) / THETA_RY * HALF_PI
                 + 3 * HALF_PI;

LUM := Y;
RG := RADIUS * sin (THETA_OPP);
YB := RADIUS * cos (THETA_OPP);
```

Appendix B

Transcript of Subject Instructions

This appendix contains the complete text of the instructions that were given to each of the inexperienced subjects prior to performing the colour matching experiment. For the most part, the instructions are shown here as they appeared on the alphanumeric display device. Where subjects were required to perform an action before paging to the next screen of instructions, the "PRESS 'f' (forward) for the next page of instructions ..." prompt did not appear until after the action had been completed. The instructions given to subjects who performed colour matching using selection systems based on the "single attribute control" mechanism are shown on pages 107 through 120. On page 121, the page of instructions that was substituted for the instructions on page 118 for subjects who used the "multiple attribute control" mechanism is shown.

University of Waterloo
Computer Graphics Laboratory

```
******************************
*                            *
*     COMPUTER ASSISTED      *
*   COLOUR MATCHING SYSTEM   *
*                            *
******************************
```

The experiment you are about to participate in is part of an investigation into human colour perception and human - computer interaction methods sponsored jointly by the Computer Graphics Laboratory of the University of Waterloo and the National Research Council of Canada.

PRESS 'f' on the keyboard to move FORWARD to the next page of instructions.

Before you begin the actual experiment, you will be led through a series of instructional procedures that will familiarize you with the computer graphics equipment and the colour matching process.

Please take your time during the instruction session. While you are not expected to be a colour matching expert before the experiment even begins, it is important that you understand how the computer graphics equipment works.

Note that the system will allow you to move BACK to previous instructions as well as FORWARD to the next set of instructions. If at any time during the instruction session or during the actual experiment you are not certain of what is required of you, please feel free to ask the experiment supervisor for assistance. Also feel free to adjust the position of any of the devices you will be using.

PRESS 'f' (forward) for the next page of instructions
or 'b' (back) for the previous page of instructions.

By now you should have been asked to take a preliminary colour vision examination. This examination was used to test your visual system for any abnormalities. If you haven't already taken such an exam, please notify the experiment supervisor.

To help in the administration of the experiment, we ask that you enter your name. Note that all personal information and the data generated by this experiment will be kept confidential.

Please type your name on the keyboard and then PRESS RETURN.
==> J. C. Blindasabat

Thank You.

PRESS 'f' (forward) for the next page of instructions
or 'b' (back) for the previous page of instructions.

Colour Matching Instructions

What you have to do is actually quite simple. On the display screen in front of you are two white squares. At the beginning of each 'colour match', the two squares will be set to different colours. The colour of the top square will stay the same for the the duration of the match. You will be able to change the colour of the bottom square. (How you do this will be described shortly). You will then manipulate the colour of the bottom square until it is the same as the colour of the top square.

Note that the two colours should be matched as closely as possible. This means that you should continue in your attempt to match the colours until you think they are the same or until it becomes EXTREMELY difficult to get the colours any closer to each other. Each match should take you about 1 to 2 minutes to complete. You will actually be allowed 3 minutes for each match.

PRESS 'f' (forward) for the next page of instructions
or 'b' (back) for the previous page of instructions.

You will be able to indicate that you have completed a match by pressing RETURN on the keyboard. After you do this, the two squares will both turn white again. The system will then wait until you press RETURN again before the next match is started. You will be asked to perform a total of 30 matches, so you may want to take a short break between matches.

PRESS 'f' (forward) for the next page of instructions
or 'b' (back) for the previous page of instructions.

Colour Manipulation Instructions

Manipulating colour, as you will probably find out, is not as simple as you may have first thought. The best way to think of colour is to picture it as some three - dimensional object such as a large box.

A box has three properties, or dimensions: width, length and height. Like a box, colour also has three properties (or dimensions). Thus, there will be three properties of colour that you will be able to control in order to change the colour of the bottom square.

If you like, you can imagine that the colour you are trying to match (ie. the colour of the top square) is a single point somewhere within the box. Your task then, is to search through the 'contents' of the box and find the 'right' point.

PRESS 'f' (forward) for the next page of instructions
or 'b' (back) for the previous page of instructions.

On the table directly in front of you is a flat, white TABLET. Sitting on the tablet and connected by a thin cord, is a device with four 'buttons' called a 'PUCK'. You will be using the tablet and puck during the experiment to alter the colour of the bottom square. The puck buttons will be referred to as 'top', 'bottom', the cord).

Each of the three properties of colour can be manipulated by HOLDING DOWN one of the buttons with your finger and sliding the puck along the surface of the tablet. For instance, if you HOLD DOWN the TOP button and move the puck to the RIGHT, you will notice that the colour of the bottom square will become BRIGHTER. Similarly, if you HOLD DOWN the TOP button and move the puck to the LEFT, the colour will become DARKER.

TRY IT and PRESS the RETURN button on the keyboard when you are convinced that this really works!

PRESS 'f' (forward) for the next page of instructions
or 'b' (back) for the previous page of instructions.

You may have noticed that the colour only changes while the button is actually held down. If you lift your finger OFF the button, so that none of the buttons are being pressed, the colour WILL NOT CHANGE when the puck is moved.

This is important because it means that if you want to increase the brightness of the colour by a large amount, you can use the following technique:

1. Move the puck to the RIGHT while holding DOWN the TOP button.
2. Move the puck back to the LEFT while keeping your finger OFF ALL BUTTONS.
3. Again move the puck to the RIGHT while holding DOWN the TOP button.
4. Move the puck back to the LEFT while keeping your finger OFF the buttons.

.
.
.

and so on.

PRESS 'f' (forward) for the next page of instructions
or 'b' (back) for the previous page of instructions.

In other words, you will be moving your hand back and forth on the tablet - holding the button down while you move in one direction and releasing it as you move in the other. Note that this means that if you want to DECREASE the brightness of the colour, you just reverse the process - hold the button down while you move to the LEFT and release it when moving back to the RIGHT.

Try using this technique. For instance, you might try making the colour of the bottom square as bright as possible. When the colour doesn't get any brighter, decrease the brightness until the colour of the square is almost black.

PRESS RETURN when you think you've got the idea:

PRESS 'f' (forward) for the next page of instructions
or 'b' (back) for the previous page of instructions.

As well as being able to change brightness, you will be able to change two other properties of colour. The best way to discover what these properties are and how the colour changes as each is altered is to actually see for yourself! Don't get too worried if you don't understand the behavior of these properties right away - you will learn this as the experiment proceeds. After all, one of the goals of this experiment is to study this initial learning period!

PRESS 'f' (forward) for the next page of instructions
or 'b' (back) for the previous page of instructions.

The second of the three colour properties you will be using is changed by HOLDING DOWN the LEFT button and moving the puck LEFT or RIGHT. The manner in which you move the puck is exactly the same as the way you used to change the brightness of the colour. The only difference is that by holding down the LEFT button instead of the TOP button, you are now changing a different colour property.

Experiment with this property for awhile and PRESS RETURN when you are ready to go on:

PRESS 'f' (forward) for the next page of instructions
or 'b' (back) for the previous page of instructions.

The third property of colour you will be using is altered by HOLDING DOWN the RIGHT button and moving the puck LEFT or RIGHT. Again, the puck movement used to manipulate this property is the same as that used to control the other properties.

Try experimenting with this property for awhile and PRESS RETURN when you are done:

PRESS 'f' (forward) for the next page of instructions
or 'b' (back) for the previous page of instructions.

By now you should feel fairly comfortable using the tablet and puck. You should also have a basic understanding of how the three colour properties are changed (even though the behavior of these properties may be unclear at this point). To summarize:

1. The first colour property (brightness) is controlled by the TOP button. This property is changed by moving the puck LEFT or RIGHT.

2. The second colour property is controlled by the LEFT button. This property is changed by moving the puck LEFT or RIGHT.

3. The third colour property is controlled by the RIGHT button. This property is changed by moving the puck LEFT or RIGHT.

PRESS 'f' (forward) for the next page of instructions
or 'b' (back) for the previous page of instructions.

Now that you have seen how each property independently affects the colour, try manipulating the three properties in combination. For example, you might try increasing the first property a little, then decreasing the second property, then increasing the third, and so on. Note that if you hold down two buttons at once, nothing will happen!

Experiment for a minute or so and then PRESS RETURN:

PRESS 'f' (forward) for the next page of instructions
or 'b' (back) for the previous page of instructions.

Subjects who performed the experiment using a colour selection system that was based on the "single attribute control" mechanism read the instructions shown on the previous 14 pages. Subjects that used selection systems based on "multiple attribute control" read exactly the same instructions except that the page of instructions shown below was substituted for the page of instructions found on page 118.

The third property of colour you will be using is also controlled by the TOP button. The difference here is that instead of moving the puck left or right to change this property, you must move it UP or DOWN. This vertical motion is analogous to the horizontal motion you have already used. That is, the colour is only affected when the button is HELD DOWN.

Since two colour properties are controlled by the same button but by moving the puck in different directions, you can now change two colour properties at the same time. This is done by sliding the puck at an angle instead of moving horizontally or vertically. For instance, if you move the puck UP and to the RIGHT at a 45 degree angle, both properties will increase at the same rate.

Experiment with this for awhile. First try to alter only the third property by moving the puck STRAIGHT up and down. Once you have an idea of how this property behaves, try moving the puck at different angles and observe the effect that these movements have.

PRESS RETURN when you are done:

PRESS 'f' to move FORWARD to next instructions
or 'b' to move BACK to previous instructions.

References

[Adage82] Adage Inc. (1982), *RDS 3000 Programming Reference Manual*, .

[Apple84] Apple Computer, Inc. (1984), QuickDraw: A Programmer's Guide, *Inside Macintosh*, Apple.

[Baecker80] Ron Baecker (1980), Towards an Effective Characterization of Graphical Interaction, *Methodology of Interaction*, D. Duce (ed.), North-Holland, 127-148.

[Becker84] Richard A. Becker and John M. Chambers (1984), *S - An Interactive Environment for Data Analysis and Graphics*, Wadsworth, Inc.

[Berk82] Toby Berk, Lee Brownston, and Arie Kaufman (1982), A New Color-Naming System for Graphics Languages, *IEEE Computer Graphics and Applications* **2**(3), May, 37-44.

[Boynton79] Robert M. Boynton (1979), *Human Color Vision*, Holt, Rinehart and Winston.

[Buxton83] William Buxton (1983), Lexical and Pragmatic Considerations of Input Structures, *Computer Graphics* **17**(1).

[CIE78] CIE (1978), Recommendations on Uniform Color Spaces, Color-Difference Equations, Psychrometric Color Terms, Bureau Central de la CIE [Supplement No. 2 of CIE Publication No. 15 (E-1.3.1) 1971].

[Card83] Stuart K. Card, Thomas P. Moran, and Allen Newell (1983), *The Psychology of Human-Computer Interfaces*, Lawrence Erlbaum Associates.

[Cowan83a] William B. Cowan (1983), An Inexpensive Scheme for Calibration of a Colour Monitor in Terms of CIE Standard Coordinates, *Computer Graphics* **17**(3), July, 315-321 [proceedings of the 1983 Siggraph Conference].

[Cowan83b] William B. Cowan and Colin Ware (1983), *Colour Perception Tutorial Notes: Siggraph '83*, Siggraph.

[Evans81] Kenneth B. Evans, Peter P. Tanner, and Marceli Wein (1981), Tablet-Based Valuators That Provide One, Two, or Three Degrees of Freedom, *Computer Graphics* **15**(3), August, 91-97 [proceedings of the 1981 Siggraph Conference].

[Fishkin83] Kenneth P. Fishkin (1983), Applying Color Science to Computer Graphics, MSc dissertation, University of California at Berkley.

[Fishkin84] Kenneth P. Fishkin and Brian A. Barsky (1984), A Family of New Algorithms for Soft Filling, *Computer Graphics* **18**(3) [proceedings of the 1984 Siggraph Conference].

[Foley82] J. D. Foley and A. Van Dam (1982), *Fundamentals of Interactive Computer Graphics*, Addison-Wesley Publications.

[GSPC79] Graphics Standards Planning Committee (1979), Status Report of the Graphics Standards Planning Committee, *Computer Graphics* **13**(3), August.

[Goetz82] Susan M. Goetz, John C. Beatty, and Deryl J. Rasquinha (1982), Colour Principles and Experience for Computer Graphics, *Graphics Interface '82*, May, 313-322.

[Guralnik80] David B. Guralnik (1980), *Webster's New World Dictionary of the American Language*, Simon and Schuster.

[Hunt75] R. W. G. Hunt (1975), *The Reproduction of Colour*, Fountain Press.

[Joblove78] George H. Joblove and Donald P. Greenberg (1978), Color Spaces for Computer Graphics, *Computer Graphics* **12**(3), August, 20-25 [proceedings of the 1978 Siggraph Conference].

[Kelly55] K. L. Kelly and D. B. Judd (1955), The ISCC-NBS Method of Designating Colors and a Dictionary of Color Names, Circular 553, National Bureau of Standards.

[Maerz50] A. Maerz and M. Paul (1950), *Dictionary of Color*, McGraw-Hill.

[Meyer80] G. Meyer and D. Greenberg (1980), Perceptual Color Spaces for Computer Graphics, *Computer Graphics* **14**(3), August, 254-261 [proceedings of the 1980 Siggraph Conference].

[Meyer83] G. W. Meyer (1983), Colorimetry and Computer Graphics, MSc dissertation, Cornell University.

[Munsell39] A. H. Munsell (1939), *A Color Notation*, Munsell Color Company.

[Ostle75] Bernard Ostle and Richard W. Mensing (1975), *Statistics in Research*, The Iowa State University Press.

[Plebon82] Darlene A. Plebon and Kellogg S. Booth (1982), Interactive Picture Creation Systems, CS-82-46, Department of Computer Science, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1.

[Singh82] Baldev Singh, John C. Beatty, Kellogg S. Booth, and Rhonda Ryman (1982), A Computerized Editor for Benesh Dance Notation, CS-82-41, Department of Computer Science, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1.

[Smith78] Alvy Ray Smith (1978), Color Gamut Transform Pairs, *Computer Graphics* **12**(3), August, 12-19 [proceedings of the 1978 Siggraph Conference].

[Smith82] Alvy Ray Smith (1982), Paint, *Tutorial: Computer Graphics*, Kellogg S. Booth (ed.), IEEE Computer Society Press, 501-515.

[Tanner83] Peter Tanner, William Cowan, and Marcelli Wein (1983), Colour Selection, Swath Brushes and Memory Architectures for Paint Systems, *Graphics Interface '83*, May, 171-190.

[Trucken81] Joan R. Truckenbrod (1981), Effective Use of Color in Computer Graphics, *Computer Graphics* **15**(3), August, 83-90 [proceedings of the 1981 Siggraph Conference].

[Wyszecki82] Gunter Wyszecki and W. S. Stiles (1982), *Color Science*, Wiley.