INFORMATION TO USERS

This manuscript has been reproduced from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send UMI a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps.

Photographs included in the original manuscript have been reproduced xerographically in this copy. Higher quality 6" x 9" black and white photographic prints are available for any photographs or illustrations appearing in this copy for an additional charge. Contact UMI directly to order.

ProQuest Information and Learning
300 North Zeeb Road, Ann Arbor, MI 48106-1346 USA
800-521-0600

UMI®

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

University of Alberta

Practice Makes Perfect or Practice Makes Permanent? The Effects of Practice on the Transfer of Prior Experience

by

Tenaha O'Reilly ©

A thesis submitted to the Faculty of Graduate Studies and Research in partial fulfillment of the requirements for the degree of Doctor of Philosophy

Department of Psychology

Edmonton, Alberta

Fall, 2001

National Library of Canada
Acquisitions and Bibliographic Services
395 Wellington Street
Ottawa ON K1A 0N4
Canada

Bibliothèque nationale du Canada
Acquisitions et services bibliographiques
395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file Votre référence

Our file Notre référence

0-612-68978-6

University of Alberta

Library Release Form

Name of Author: Tenaha O'Reilly

Title of Thesis: Practice Makes Perfect or Practice Makes Permanent? The Effects of Practice on the Transfer of Prior Experience

Degree: Doctor of Philosophy

Year this Degree Granted: 2001

Permission is hereby granted to the University of Alberta Library to reproduce single copies of this thesis and to lend or sell such copies for private, scholarly or scientific research purposes only.

The author reserves other publication rights, and neither the thesis nor extensive extracts from it may be printed or otherwise be reproduced without the author's written permission.

Tenaha O'Reilly
170 Lincoln Road
Grand Falls - Windsor, N.F.
A2A 1P6

Date: Aug 31/01

University of Alberta

Faculty of Graduate Studies and Research

The undersigned certify that they have read, and recommend to the Faculty of Graduate Studies and Research for acceptance, a thesis entitled Practice Makes Perfect or Practice Makes Permanent? The Effects of Practice on the Transfer of Prior Experience submitted by Tenaha O'Reilly in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

Peter Dixon, Ph.D.

Alinda Friedman, Ph.D.

Norman Brown, Ph.D.

Craig Montgomerie, Ph.D.

Richard Catrambone, Ph.D.

Date: Aug 28/01

Abstract

The purpose of the present research was to investigate the influence of practice and multiple examples on a learner's ability to acquire a novel task. Participants first learned how to operate several electronic weather devices and then attempted to transfer their knowledge to a number of related, but superficially different, weather devices. Superficial similarity was manipulated in terms of the layout of the interface (interface layout) and the choice of control name (control names). In some conditions, participants were given practice with operating a single example of a device; in other conditions, participants were given practice with operating two related devices (multiple examples). According to a generalized view of practice, practice with a particular device should enable the user to develop a flexible representation that can aid in the acquisition of superficially dissimilar devices. In contrast, the specialized view of practice asserts that practice enables the user to optimally execute only the practiced task, and transfer to related, but different tasks is limited. When participants practiced a task with a single example of the device, the effects of practice were specialized. However, when participants were provided with multiple examples of a device, the effects of practice were general. These results suggest that the type of practice is critical in determining how effectively the user can transfer their prior experience to a superficially dissimilar device.

Dedication

This dissertation is dedicated to the many people who have enriched and inspired my life.

First of all, I would like to thank Peter Dixon for shining the "light of reason" upon me. Without your patience and guidance, this dissertation would not have been possible. You taught me more than you will ever know.

I would like to thank Gay Bisanz, Alinda Friedman, Norman Brown, Craig Montogmerie, and Richard Catrambone for your support and comments on the dissertation.

I thank Christa Philpott for her support and encouragement throughout this endeavor. I don't know what I would do without you.

Special thanks to the many friends who stood by me during times of laughter and despair. In particular, Colin Hillyard, Forrest Thompson, Kevin Philpott, and Richard Fowler.

But most important of all, I thank the two most important and greatest people in the world, my Mother and Father. Your unconditional love and support made me what I am today. Words can't express my love and gratitude towards you. I hope I made you proud.

One final note: Always take time to support, listen and encourage others- sometimes you never realize how much of an impact you can have on a person's life.

Table of Contents

Chapter 1

Chapter 2

Chapter 3

Chapter 4

Chapter 5

Chapter 6

List of Tables

List of Figures

Chapter 1

Introduction

The question of how old knowledge can be applied to novel situations has been one of the most scrutinized areas of both cognitive and educational psychology. The transfer of training from one task to another has been investigated by researchers in domains such as text editing (Singley & Anderson, 1988); computer programming (Harvey & Anderson, 1996); concept formation (Schroth, 1995); probability (Catrambone, 1995; Ross & Kilbane, 1997); algebra (Novick, 1988); insight learning (Perfetto, Bransford & Franks, 1983) and in the operation of devices (Dixon & Gabrys, 1991; Dixon, Zimmerman & Neary, 1997). Transfer has also been examined in more general areas such as problem solving (Catrambone & Holyoak, 1989; Gick & Holyoak, 1980) and expert performance (Chase & Ericsson, 1982; Ericsson & Polson, 1988; Novick, 1988). On the one hand, spontaneous transfer between diverse tasks has been demonstrated when learners were unaware of the influence of prior training on their ability to solve a novel problem (Schunn & Dunbar, 1997). On the other hand, other work has failed to show positive transfer between tasks even when the information from one problem blatantly suggested the solution for another problem (Perfetto et al. 1983; also see Weisberg, DiCamillo & Phillips,1978, Experiment 2).

Despite the seemingly intractable nature of knowledge transfer, recent research has lead to some insightful theories of how people are able to utilize their prior knowledge in the context of novel problems. These include production models (e.g., Kitajima & Polson, 1997; Singly & Anderson, 1988), models of analogical reasoning such as the structure mapping model (e.g., Forbus & Gentner, 1994) and pragmatic schema model (Analogical mapping by constraint satisfaction, ARCS) (e.g., Wharton, Holyoak, Downing, Lange, Wickens & Melz, 1994) and exemplar models such as Hintzman's (1986) trace model and Ross' reminding theory (1984,1987,1989). While each theory has its own relative merits and weaknesses, all models have contributed to our understanding of knowledge transfer. One of the goals of the proposed research is to determine whether the learning model of O'Reilly and Dixon (2001a,b) can account for the effects of practice on the transfer of prior device experience.

Other efforts have been focused on ways to improve transfer. Some of these include the provision of a hint (Gick & Holyoak, 1980; Weisberg et al. 1978); transfer appropriate processing (Needham & Beg, 1991); labeling (Catrambone, 1996); prompted review (Trudel & Payne, 1996); providing structure diagrams (Van der Veer, 1989); structural training (Munley & Patrick, 1997); comparing multiple examples (Gick & Holyoak, 1983); self-explanation (Chi, Bassok, Lewis, Reimann, & Glaser, 1989; Howie & Vincite, 1998); and exploration (Kamouri, Kamouri & Smith, 1986). The effectiveness of these interventions often depends on the type of material to be learned, but attempts to improve the transfer of training have been promising. In the present article, one such attempt to improve transfer is explored: practice.

Practice is a very simple learning strategy that can be used to acquire new information. In fact, repetition is often used as the default strategy for learning verbal material (Gardner, 1990). The simplicity and common usage of this strategy in the domain of verbal learning makes it a potentially good candidate for increasing transfer in procedural domains. While other strategies for improving transfer have been effective, they are more effortful and as a result few of them are being used in the classroom. This is especially the case for metacognitive interventions (Moley, Hart, Leal, Santulli, Rao, Johnson and Hamilton, 1992). Furthermore, it is also useful to examine the effects of practice on transfer for reasons other than strategic intervention. For example, suppose that a factory has been using a particular device interface for manufacturing their product for several years. During that time it is expected that the workers would have become very proficient at operating the existing interface. However, with increasing demand and ever increasing advances in technology, the company will be forced to modernize their system over time. The modernization process could lead to the development of a new interface that carries out similar tasks as the old one, but the interface might be arranged differently and the over all appearance of the device may be different as well. An interesting question is whether the workers with extensive experience on the old system will be able to transfer their prior experience to the new system without any difficulty. The proposed research is designed to provide some insight into this question.

Research indicates that practice generally improves performance on tasks that are very similar to the practiced task, but transfer to less related tasks is limited (e.g., Anderson, 1987; Harvey and Anderson, 1996; Jolicoeur, 1985; Logan, 1988; Singley & Anderson, 1989). This robust finding is hereafter called the *specialized view* of practice. However, a small body of literature in the area of multiple examples (Catrambone & Holyoak, 1989; Gick & Holyoak, 1983; Jolicoeur & Milliken, 1989 Experiment 2; Ross & Kennedy, 1990), expertise (Chi, Feltovich & Glaser,1981; Novick, 1988) and device operation (Bibby & Payne, 1996) suggests that practice can facilitate the transfer of prior knowledge to related problems and tasks. This finding is herein referred to as the *generalized view* of practice. The purpose of the proposed work is to determine whether the generalized or the specialized view of practice can account for the transfer of prior device experiences.

For example, a proponent of the generalized view might predict that practice would allow a user to extend his or her knowledge of a Sony VCR to help them use other VCRs (e.g., Pioneer), and possibly, even a DVD player. During practice the user might develop a deep, flexible representation of the task, and accordingly, they should be able to extend this representation to related situations. In contrast, if the effects of practice were specialized, then one might expect that practice with a Sony VCR would help the user to carry out the tasks and procedures more efficiently on the Sony model. However, this experience would not be expected to help the user to use a Pioneer VCR because the procedures for carrying out the tasks on the Pioneer are slightly different and the system has a different interface. On this account, one develops a superficial representation of the task that is specific to the device at hand. More generally, a proponent of the generalized view would assert that practice helps the user extend their specific experience to related situations, while the specialized view asserts that practice facilitates performance on the practiced task only.

Research on the initial stages of learning has indicated that novices often have difficultly transferring their knowledge of a previously encountered device to a superficially different device (O'Reilly & Dixon, 2001b; Schumacher & Gentner, 1988). This is especially true when the learners are not provided with a hint that relates the training and transfer devices (O'Reilly & Dixon, 2001a,b). The goal of the proposed experiments is to determine whether practice can override the normal tendency for novices to be grounded in the superficial features of a given device. Furthermore, I am interested in whether the device-learning model proposed by O'Reilly and Dixon (2001a,b) can be extended to account for the effect of practice on the transfer of device knowledge.

I chose to examine transfer in the domain of device operation for several reasons. First, the sheer volume and wide spread usage of electronic devices in the real world provide a representative testing ground for examining transfer in ecological settings. With new and more modern versions of devices appearing on the market daily, the amount of opportunities for the transfer of device experience are phenomenal. Furthermore, by examining the factors that are responsible for the transfer of device knowledge, researchers can develop more effective interventions for learning how to use and apply device knowledge. Such strategic interventions can help combat the tendency to avoid using electronic devices that would otherwise help people to execute their tasks more efficiently (Nickerson, 1999).

The first part of the paper provides a history of the early work on transfer. Then, the literature on learning at the novice level is reviewed. This research includes learning from one example and learning from multiple examples. The research on learning from one example is provided to introduce the reader to one of the most ubiquitous findings in the transfer literature: the superficial similarity effect. This line of work suggests that novices are often influenced by irrelevant aspects of a problem rather than the more important structural aspects (e.g., Holyoak, & Kuh, 1987). The goal of the proposed work is to discover ways to reduce the effects of superficial similarity. One such attempt to reduce these effects is to provide learners with multiple examples. The work on multiple examples suggests that the effect of superficial similarity can be reduced if then learner solves multiple examples of a principle (e.g., Catrambone & Holyoak, 1989). Next, another attempt to reduce the effect of superficial similarity is explored: practice. It is argued that providing multiple examples is a special form of practice. Subsequently, I review the work of O'Reilly and Dixon (2001a,b) on the transfer of device knowledge. This section is included to provide a background of my work and to introduce the learning model proposed by O'Reilly & Dixon (2001a,b). This section also provides a context for the present experiments. Finally, I describe the five experiments conducted to determine whether the effects of practice are specialized or general and whether the learning model proposed by O'Reilly and Dixon (2001a,b) can account for the effects of practice and varied training (multiple examples) on the transfer of device experience.

Historical Perspectives on The Transfer Of Training

For more than 100 years scientists have been trying to understand how prior experience can be utilized in the acquisition of novel problems. In the early nineteenth century, scholars believed that transfer across a variety of domains was possible through a practice called the "Doctrine of Formal Discipline" (Angell, 1908). According to this view, the mind is like a set of muscles that require constant exercise in order to perform at an optimal level. The educators of the day alleged that studying arduous subjects such as Greek, Latin, and geometry would function to "discipline" the mind. It was believed that the practice of such challenging subjects would lead to the facilitation of general faculties such as attention, discrimination, observation, and reasoning. The content or focus of study was thought to be of little importance; however, the amount of effort required in performing the activity was paramount. Transfer was considered to be so broad and general that it could take place across domains that shared little subject matter. For instance, transfer would be expected between riding a bike and sailing a ship because both involve the faculty of attention.

In a series of papers, Thorndike and Woodworth argued against the doctrine of formal discipline (Thorndike & Woodworth, 1901a, 1901b, 1901c). They proposed that the general faculties held by the proponents of formal discipline were essentially incorrect. According to Thorndike and Woodworth, the scope of transfer was much more specific than was previously believed by the advocates of formal discipline. Thorndike and Woodworth proposed their theory of "Identical Elements" which stated that the critical determinant of transfer is the amount of similarity between the features of the tasks. In support of their theory, Thorndike and Woodworth found that there was only modest transfer between the estimation of the area of a rectangle and the estimation of the area of a triangle (Thorndike& Woodworth, 1901a; 1901b). Other work showed that practice with locating French verbs in a text did not transfer to a task that required locating English verbs (Thorndike& Woodworth, 1901c). Even more striking was the finding that training in locating the letter *e* in text did not transfer to a task that required the location of other letters such as *s* and *p*, and *i* and *t* (Thorndike and Woodworth, 1901c).

In essence, Thorndike and Woodworth believed that transfer should improve as the function of the number of identical elements common to both tasks increased. In other words, transfer is due to stimulus-response generalization and to the formation of "habits" over time. However, in contrast to popular belief, Thorndike's view of transfer was not entirely behaviourist in nature; in fact, his ideas contained a cognitive element: selection (Cox, 1997). For instance, Thorndike and Woodworth proposed that learners have the ability to be selective: They have control over the essential features that they choose to extract from the tasks. In addition, the researchers proposed that the relations amongst these common features were formed and strengthened by practice and formed what Thorndike and Woodworth called "habits" (Thorndike& Woodworth, 1901a). In sum, Thorndike and Woodworth's theory of identical elements was one of the first systematic and scientific attempts to study the issue of transfer.

However, not everyone welcomed this atomic view of transfer, especially the proponents of the Gestalt movement. According to the Gestalt psychologists, transfer was much more holistic than the behaviourists like Thorndike and Woodworth believed. Instead of breaking down a task into its elementary components, the Gestalt theorists believed that a problem should be viewed as a whole. After World War I, the famous Gestalt psychologist Wolfgang Koehler (see Koehler, 1927) was trapped on a small island of Tenerife in the Canary Islands. During that time Koehler began to study a group of chimpanzees. From his observation on the island and other work, Koehler noticed that his chimpanzees would struggle with a problem for long time, then suddenly change their behaviour and effectively solve the problem. Moreover, Koehler found that the apes could transfer this knowledge to related tasks with no difficulty. According to Koehler, this type of learning was different from stimulus-response learning; this type of learning was what he called insight learning. The apes could not have solved the problem unless they could abstract the meaningful relations from the environment. He also argued that the meaningful relations could not be accessed if the environment could not be perceived as a whole. In other words, the holistic perception of the environment affords the necessary relations to solve the problem, and when these relations are noticed, insight learning occurs.

This holistic view of learning was applied to the classroom. According to Wertheimer (1959), teachers should avoid direct instruction because it creates habits. (This is in direct contrast to Thorndike's view.) These habits could in turn cloud the mind of the learner and effectively rule out other possible solutions to the problem. Instead, he advocated teaching by discovery in which the learner is led to grasp the principle through guided application of the concept in terms of real life examples. By relating the principle to real world applications, the student effectively learns when the principle will and will not work. In sum, the goal is to practice with a wide variety of stimuli in order to avoid habits that restrict a person to only consider a particular response. In this way, the principle becomes context independent and a part of the "structural whole."

In summary, nineteenth century educators suggested that transfer was mediated by a set of general mental faculties. In direct contrast, Thorndike and Woodworth proposed that transfer was much more specific than what the proponents of the doctrine of formal discipline asserted. According to Thorndike and Woodworth, transfer is a function of the number of identical elements that are common to both tasks. Transfer is facilitated by strengthening the relations between the elements of the tasks through simple practice. In reaction to the behaviourist tradition, the Gestaltists believed that the issue of transfer should be investigated in the context of insight learning; problems should not be broken down into their elementary components, but rather they should be studied as functional wholes. According to the Gestaltists, providing wide and varied practice was believed to be the most effective way of increasing the amount of transfer to novel domains.

The next section describes some modern research on the transfer of training. First, the research on learning at the novice levels and the work on learning from a single-example are described. It will become evident that much of the current work on transfer contains ideas that are rooted in the historical accounts of transfer.

Transfer at the Novice Level: The Influence of Superficial Similarity.

The following section reviews the research on the transfer of training at the novice level. Much of this work has focused on how people learn a single principle (e.g., Newton's Second Law, $F=ma$) from studying a single example (VanLehn, 1996). In recent years, investigators have become increasingly interested in how people learn from examples (Anderson, Fincham, & Douglass, 1997; Catrambone, 1996; Ross & Killbane, 1997; VanLehn, 1998). It is often thought that example studying plays a critical part in the development of most cognitive skills (Chi, et al., 1989). In fact, for domain novices, learning from examples is the preferred method of learning compared to other forms of instruction (Chi, et al., 1989; LeFevre & Dixon, 1986; Pirolli & Anderson, 1985). Examples can help participants learn a principle by illustrating how and when a principle should be applied (Reeves, & Weisberg, 1994), and research has shown that principles are learned more easily when they are incorporated in the context of an example (Cheng, Holyoak, Nisbett & Oliver, 1986; Fong, Krantz & Nisbett, 1986).

In a typical experiment of this sort, participants learn a principle by studying an example that illustrates the application of the principle. Then, in a later part of the experiment, the learner is asked to solve another task or problem that is different, but related to the training problem. The conclusion that can be drawn from this line of research is that domain novices often have great difficulty in transferring their knowledge to novel problems when the problems share a dissimilar surface structure (e.g., Gentner,

& Markman, 1997; Heydenbluth, & Hesse, 1996; Holyoak, & Kuh, 1987; Novick, 1988; Ross, 1984, 1987,1989).

In cognitive psychology, researchers often make a distinction between a problem's "deep" and "surface" structure (Chi, Feltovich, & Glaser, 1981). The surface structure refers to the setting, events, and objects described in the problem, while the deep structure refers to the principles, laws, concepts, or equations involved in solving the problem (Blessing & Ross, 1996). Only the deep aspects of the problem are necessary for a learner to successfully solve a problem. The surface or superficial aspects are not related to the solution of the problem but are included in the problem for other reasons. For instance, consider the following physics problem:

> Jeff loves to fly the Cessna 170 aircraft. At an altitude of 2500 feet, the Cessna 170 generates 123 horsepower at 2700 RPM. Under these conditions, the 170 can fly at a range of 455 miles with a lean mixture of fuel. According to the logbook, on June 27, 1985, Jeff traveled a distance of 200 miles in 2 hours. Assuming that Jeff traveled at a constant rate, what was his speed?

In this example, the deep structure of the problem would include the equation needed for the solution (velocity = distance / time) and the relevant numbers (200 miles and 2 hours) used in calculating a solution. On the other hand, the superficial features of the problem would include the name of the plane (Cessna 170), its specifications (e.g., 123 horsepower, 455 mile range), the pilot's identity (Jeff), and the date of the flight (June 27, 1985). None of this information is relevant for calculating the solution, and ideally, it should be unrelated to a learner's ability to reach the correct solution.

However, a large body of evidence suggests that domain novices are often influenced by the superficial characteristics of a problem (e.g., Gentner, & Markman, 1997; Heydenbluth, & Hesse, 1996; Holyoak, & Koh, 1987; Novick, 1988; Ross, 1984, 1987,1989). For example, Chi et al. (1981) asked both novice and expert physics students to categorize various physics problems. Participants were asked to sort a set of problems into groups based on the similarities of their solution. The experimenters found that the physics novices categorized the problems based on the superficial features, while the physics experts sorted the problems in terms of their deep or structural aspects. For instance, the experts were more likely to classify the problems based on their underlying principle (e.g., Newton's Second Law, $F= ma$), while novices tended to sort them by the objects (springs, inclined planes) mentioned in the problem. These findings indicate that a learner's level of expertise greatly influences how he or she categorizes and represents the problems in a domain. In particular, superficial similarity has a larger influence on novices' ability to classify problems than it does for experts.

However, the effect of superficial similarity is not limited to a novice's ability to merely categorize a problem; it also has an effect on one's ability to access and use prior experience with similar problems (Heydenbluth & Hess, 1996; Holyoak & Koh 1987; Ross, 1987, 1989). For example, Ross (1989) found that math novices were more likely to provide the correct solution for a probability problem if the test problem contained the same story line and objects as the problem that they were trained on. For instance, if a participant was trained on a problem about golf (story line) and golfers (objects) they were more likely to propose the correct solution of a related problem when the test problem was about golf and golfers than when it concerned dancing and dancers, despite the fact that the two types of test problems were formally identical. Ross proposed that the superficial similarity in the problem reminded the learners that the training and transfer tasks were related (e.g., Ross, 1984). With this information available, learners could transfer their knowledge of one problem onto another.

Other work has indicated that the transfer of prior knowledge is poor when the training and transfer tasks are superficially dissimilar. In this manner, superficial *dis*similarity can have a negative effect on a domain novice's ability to transfer their knowledge to a structurally related task. For example, even when learners have the knowledge needed to form a successful correspondence between the training and transfer tasks, they often fail to use the knowledge to help them solve the problem (Catrambone & Holyoak, 1989; Stein, Way, Benningfield, & Hedgecough, 1986; Weisberg, DiCamillo, & Phillips, 1978).

However, providing learners with a hint about how to approach the problem can have a large effect on transfer (e.g., Ross & Kennedy, 1990; Weisberg et al. 1978). A clear example of the role of providing a hint can be found in the research of Gick and Holyoak (1980). Participants were first given Dunker's ray problem: A patient has developed a malignant tumor, and without treatment the patient will die. The tumor can be destroyed by using high levels of rays, but at these levels the healthy tissue will also die. The critical solution is to focus low levels of radiation on the tumor from different sides. This was preceded by an analogous problem in a different domain: A general wants to take over heavily guarded fortress surrounded by mines. Because only a few soldiers can cross the minefield at a time from any direction, the attack must occur in small groups from all sides. Only about 20% of the participants spontaneously used the general solution to solve the ray problem. However, when participants were explicitly told the relationship between the two tasks, they were much more likely to generate the analogous solution. In this case about 92% of the participants were able to apply the general solution to the ray problem. Thus, when participants are explicitly asked to adopt a strategy of comparing the two problems, they notice the correspondences and find an analogous solution.

In sum, domain novices often prefer to learn a principle or skill by studying examples (Chi, et al., 1989; LeFevre & Dixon, 1986; Pirolli & Anderson, 1985). However, example studying can be problematic as a method of learning for domain novices. Novices not only learn the important or deep aspects contained in an example, but they also pick up on the irrelevant or superficial aspects. The superficial nature of a problem can influence how a domain novice represents (Chi et al. 1981) and applies (Ross, 1987, 1989) a problem solution. In many cases, this superficial representation and application of an example solution can hinder the transfer of prior knowledge (Heydenbluth & Hesse, 1996; Holyoak & Koh, 1987; Ross, 1989; Schumacher & Gentner, 1988). Moreover, domain novices often require a hint in order for them to successfully transfer their knowledge of one task to another (Gick & Holyoak, 1980; Weisberg et al. 1978). Thus, the research on learning from a single-example seems to suggest that the novice's knowledge of a domain is relatively superficial: Spontaneous transfer to superficially dissimilar tasks is poor, and transfer is usually less when the task is superficially dissimilar even when the learners are provided with a hint (e.g., Ross, 1989).

The research on learning from a single example leads to a conclusion similar to that reached by Thorndike and Woodworth a hundred years earlier: Transfer is dependent upon the similarities (identical elements) that exist between the training and transfer task. A number of researchers have concluded that some of the modern theories of psychology are merely "old theories in modern dress" (Anastasi, McKeachie, Smith, & Cummings, 1992). Some argue that this is especially the case with respect to the behaviourist approach to transfer (e.g., identical elements theory) and modern production models (Butterfield, Slocum & Nelson, 1993). For example, one could argue that the condition and action components found in the production models is simply a substitute for the stimulus and response terms used by the behaviourists.

Learning From Multiple Examples: Generalization?

The research on learning a principle from a single example suggests that learners often display poor transfer to novel problems when the surface structure is different from that on the training example. However, is it possible for a learner to reduce the influence of superficial similarity by training participants on a variety of problems that illustrate the same principle? According to this view, learners can develop generalizations across different problem types by extracting the essential structure (i.e., the principle) from each instance. Through multiple examples, learners could develop a problem schema that involves deleting the differences between the problems and preserving their commonalities (Gick & Holyoak, 1983). In other words, providing a dissimilar training regime might make the task more difficult to learn, but transfer might be more flexible afterwards (Gick & Holyoak, 1983). In essence, this view is reminiscent of the insight learning and pedagogy of the Gestalt tradition (e.g.,Wertheimer, 1959). According to Wertheimer, providing varied practice should promote an increase in the transfer of prior knowledge. It could be the case that domain novices are affected by the superficial features of the problem simply because their limited experience does not afford the opportunity to spontaneously extract the critical structural elements from the problem.

Recent work seems to support this generalization view (Catrambone, & Holyoak, 1989; Gick & Holyoak, 1983; Jolicoeur & Milliken, 1989 Experiment 2; Ross & Kennedy, 1990). For, example, Gick and Holyoak (1983) found that participants were more likely to solve a test problem if they were explicitly asked to compare two problems that illustrated the same principle than if they were provided with a single-example that illustrated the principle. Without a hint, 45% of the participants who were given multiple examples solved the problem, in contrast to 21% of the participants who were given one example. In addition, the researchers found that the participants who wrote better descriptions of the structural similarities between the examples (schema quality) were more likely to arrive at the correct the solution: among participants who wrote good schemas, 91% produced the solution without a hint; among those who produced an intermediate-quality schema, 40% produced a solution; and among participants who wrote poor schemas, the number was only 30%. The authors attributed successful transfer to the formation of a schema that allowed the application of general rules to other problems. However, subsequent work failed to show any advantage of providing multiple examples when the context of the training and transfer tasks were different (Spencer & Weisberg, 1986).

In other work, Catrambone and Holyoak (1989) rebutted the findings of Spencer and Weisberg, (1986) on the grounds that most of their participants wrote poor schema descriptions of the examples. Catrambone and Holyoak (1989) replicated Gick and Holyoak's (1983) findings and concluded that the important determinant of transfer is not the number of examples per se but rather the quality of the schemas produced by the participants. As in the Gick and Holyoak (1983) study, if the quality of the schema was poor, then transfer to novel problems was also poor. However, Catrambone and Holyoak (1989) also found that transfer was dependent upon whether the learners explicitly compared the examples rather than if they summarized them. About 47% of the participants who were instructed to compare the examples solved the transfer problem as opposed to only 16% of the participants who were asked to merely summarize the examples (see Experiment 1). While these findings seem promising, Catrambone and Holyoak (1989) found that transfer without a hint was very low after a delay of a week (see Experiment 2). However, in Experiment 4, the researchers found that pre-hint transfer was increased when the participants were given extensive comparison instructions that allowed the participants to form a more effective problem-solving schema. They concluded that processing multiple examples is not sufficient enough to promote transfer. Instead, the learner must be directed to focus on the relevant aspects of the problem. In sum, the experimenters concluded that the provision of multiple examples, direct comparison instructions, and problem solving experience contributes to the acquisition of a generalized schema. The combination of all of these factors allows more flexible transfer. On the one hand, the research by Catrambone and Holyoak (1989) is encouraging because it supports the notion that the provision of multiple examples can cause learners to develop generalizations that allows for more flexible transfer. On the other hand, the participants in their study were only able to generalize their knowledge when they explicitly compared the examples in the appropriate manner (i.e., schema quality).

In contrast, other work has indicated that generalization can occur under less restrictive conditions. For example, Ross and Kennedy (1990) provided participants with several examples of a probability principle, and then asked the participants to solve another problem that was related to the training examples. When the experimenter informed the learners that the training examples were related to one another, the participants were able to successfully transfer their knowledge to a superficially different transfer task. The encouraging finding of the Ross and Kennedy study is that the learners only required the aid of a hint in the training tasks; the participants were successfully able to transfer their knowledge to a superficially dissimilar task without a hint that related the training and transfer tasks. The experimenters concluded that cueing (relating) the training examples leads to generalization and can be used to overcome the problems associated with transfer of superficially dissimilar examples. However, they did indicate that the comparison process in their task was not automatic.

In contrast, other work indicates that automatic schema abstraction (i.e., without explicit instruction) can occur when the learner has been provided four or more examples (Reeves & Weisberg, 1990), or when the participant's prior knowledge is high (Ahn, Brewer, & Mooney, 1992). However, this "automatic schema abstraction" is more of an exception rather than the norm (Reeves & Weisberg, 1994). Finally, generalizations have been found in domains other than problem solving. For instance, providing participants with varied practice can increase the amount of transfer on a picture-naming task relative to a group of participants who are not given such varied training (Jolicoeur & Milliken, 1989).

In sum, the research on learning from multiple examples of a single principle indicates that learners are capable of developing generalizations or schemas that facilitate the transfer of prior knowledge to superficially dissimilar domains. The experiments indicate that schema abstraction or generalization is only possible when participants are required to compare multiple examples (Ross & Kennedy, 1990), and automatic schema abstraction is infrequent. In addition, sometimes transfer is also dependent upon the quality of these comparisons (Catrambone & Holyoak, 1989). Almost 75 years later, the Gestalt approach to problem solving has stood the test of time: The provision of varied practice can be an effective way to promote the transfer of prior knowledge.

The Influence of Practice on Transfer: The Case of Specificity?

This section reviews the evidence for the effects of practice on the transfer of prior experience. The bulk of this research indicates that the effects of practice are specialized and that transfer to related tasks is limited (e.g., Anderson, 1987; Logan, 1988; Jolicoeur, 1985). However, there are some hints that that practice might facilitate the transfer of prior experience to dissimilar tasks (e.g., Anderson, Fincham, Douglass 1997; Bibby & Payne, 1996). The section begins by providing a brief overview of the effects of practice on performance, and then the work on the influence of practice on transfer is discussed.

Practice and Performance.

In verbal learning tasks, repetition is often used as a default strategy when average students attempt to learn novel material (Gardner, 1990). However, a repetition-based strategy is often an ineffective method of learning when compared to more potent techniques such as elaboration (Willoughby, Wood, & Kahn, 1994) and self-explanation (Chi, De Leeuw, Chiu, & LaVancher, 1994). In direct contrast, performance of activities or procedures can greatly benefit from practice (e.g., Crossman, 1959; Kolers, 1976). When procedural tasks are practiced over a period of time, two robust findings occur. First, there is a general speed up in the time that it takes to perform the activity, and second, there is generally an increase in the accuracy or the appropriateness of the actions for the task (VanLehn, 1996).

The speedup in task execution is often referred to as the power law of learning (Anderson, 1995). That is, the time required to complete a given activity decreases according to a negative power function of the form: $T = A+BN^{-c}$. Where T is the amount of time to complete the task; A is the asymptote, B is a scale factor (or the amount of time that can be improved by practice), N is the number of practice trials, and c is an exponent that determines the rate of speedup with practice. (However, recent evidence suggests that the effects of practice may be best modeled by an exponential function; Heathcote, Brown & Mewhort, 2000.) For instance, Crossman (1959) examined the ability of workers to produce cigars over a 10-year period. His observations revealed that there was large decrease in the time required producing a cigar when the workers were given one year of practice. However, this decrease in cycle time diminished over time: Within 5 years of practice, the initial speedup in cycle time had reached a relative asymptotic level, and further practice did not result in a significant decrease in production time. Similarly, the time to solve geometry proofs decreases as the number of problems solved increased. However, as with the cigar study, the speed up followed a negative power law; thus, the reduction in solving time associated with each problem solution decreased over the number of problems solved (Neves & Anderson, 1981).

According to Fitts (1964), the beneficial effects of practice on cognitive and motor skills generally occur in three stages. In the *early* stage, learners are focused on understanding all of the domain-relevant knowledge of the task. Learners read the relevant information about the topic, and they also learn the set of actions to execute the task. In the *intermediate* stage, the learner begins to perform the task or skill with reference to the declarative knowledge of the domain. In this stage, the learner's performance is slow because they either have not acquired all of the domain knowledge or have inefficient methods of retrieving and applying the knowledge. Students in the intermediate level are prone to errors and misunderstandings of the domain. As a result, learners may often refer to the acquired knowledge base or other sources when performing the activity. The goal of this stage is to remove the flaws and apply more effective and faster ways of carrying out the task (e.g., heuristics). In the *late* stage, learners improve in speed and accuracy, even though their knowledge of the domain remains relatively unchanged.

From this brief discussion, it seems clear the effects of practice on task acquisition are beneficial; that is, there is a clear speedup and increase in accuracy of the performed skill over time. However, it might be possible that practice can also harm performance. Practice could lead the learner to confine or bias their course of action towards a mode of processing normally associated with the practiced activity. In other words, increased practice might lead to poor transfer and generalization.

Research that Supports the "Specialized View" of Practice.

This section reviews the research that supports the notion that the benefits of practice are specialized and that transfer to related tasks is limited (e.g., Anderson, 1987; Harvey and Anderson, 1996; Jolicoeur, 1985; Logan, 1988; Singley & Anderson, 1989). For instance, Jolicoeur (1985) asked participants to name a set of misoriented pictures as quickly as possible. Initially, the participants were much slower at naming the misoriented pictures than they were at naming the upright pictures. This is often referred to as the orientation effect: The time to name misoriented pictures increases as the distance increases between the top of an object and its usual upright position (Jolicoeur

& Landau, 1984; Kolers & Perkins, 1969). However, with several blocks of practice, the participants were much faster at naming the misoriented pictures than they were initially. In other words, practice helped reduce the orientation effect.

In another experiment, Jolicoeur (1985, Experiment 3) had participants name one set of rotated pictures for five blocks of practice, and a different set of stimuli on the sixth block. As in the previous experiments, the orientation effect decreased with practice during the first five blocks of trials. However, on the sixth block, the orientation effect reappeared when the new stimuli were introduced. Thus, the reduction of the orientation effect was not due to some general ability to name disoriented pictures but to an increased ability to name disoriented versions of only the pictures that had been practiced. A similar finding led Jolicoeur and Milliken (1989) to conclude that participants do not automatically acquire a general ability to name misoriented pictures because transfer to unlearned orientations is low.

In another domain, Logan (1988) found that the time to make a lexical decision decreased as a function of practice. In the training phase of the study, participants were required to indicate whether the presented item was a word or not. In some cases the item was a word and in other cases the presented item was a non-word. In the transfer phase of the study, participants were again required to make a lexical decision on both the words and non-words that they saw in training and an additional set or words and non-words that they had never seen before. Logan found that the time to name both the old words and the non-words decreased in the transfer phase of the experiment. However, the decision times for the new words and new non-words did not decrease. In other words, the facilitating effects of practice did not generalize to new items.

In the domain of computer programming, research has also shown that the effects of practice on the transfer of prior experience are specialized. For example, Harvey and Anderson (1996) examined novices' ability to transfer between two different programming languages: Lisp and Prolog. The programming languages are conceptually similar but syntactically different. That is, both languages have similar functions for performing certain operations, but the steps and commands used to carry out the tasks are different. For example the command to extract the first item in a list is achieved by using the *tail operator* in prolog, and the *car* operator in Lisp. In Experiment 1, participants read some instructions on how to program in Prolog, and then they attempted to solve some problems using the language. During the transfer phase, the participants read instructions on how to program in Lisp and then attempted to solve problems in Lisp. The experimenters measured the amount of time that participants took to read the pages in the manuals. Harvey and Anderson found that the participants who read the Prolog instructions were much faster at reading the Lisp instructions than a control group who did not receive any training with Prolog. However, this effect did not generalize to all pages in the Lisp instructions. Transfer was highest for pages that were identical to the Prolog pages and lowest for pages that were particular to one language. Furthermore, there was no difference in the time to write the code, suggesting that transfer did not generalize to the dissimilar syntactic and procedural aspects of the tasks. Transfer was only evident in the common conceptual aspects of the languages; there was no transfer of syntactic knowledge. Experiment 2 was similar to Experiment 1 with the exception that the participants were experienced programming in languages other than Prolog and Lisp. In addition, the amount of practice on the training task was doubled. Despite the increased amount of practice on the training tasks, the findings were essentially the same as in the first experiment: Transfer occurred across the conceptual aspects of the task, but not the procedural aspects.

The specificity of practice has even been demonstrated in real-world contexts. A real-world example of how extensive experience does not transfer to related tasks comes from a study of Bazillion street vendors (Carraher, Carraher & Schliemann, 1985). In many parts of Brazil, it is common for the sons and daughters of street vendors to help their parents with the family business. In this type of work, the children often have to perform a large number of mathematical operations. For instance, they need to be able to calculate the total cost of a number of items and how much change they should give the customer. The children carry out these operations without a calculator, and many of them have little or no schooling. Surprisingly, the children are quite good at calculating the correct solutions to the problems that they encounter at the market (e.g., How much is a coconut? How much would 10 be?). For instance, Carraher, Carraher & Schliemann (1985) found that the street vendors were able to correctly calculate 98% of the problems that were presented to them at the markets. Sometime later, the experimenters asked the vendors to solve some story problems involving the same numbers and operations that they had used in the vending context. In this case, the children correctly solved 74% of the problems. In the final phase of the experiment, the children were asked to solve some problems in the form of mathematical equations (e.g., 2*10=?). The numbers and the operations used in the mathematical equations were the same as those that were asked in the vending context. In this case, the children only solved 37% of the problems. The street vendor's extensive and successful experience with calculating mathematical problems in the markets did not transfer well to the same problems presented in a different context. In this case, the street vendor's skills were very specialized, and transfer to a highly related task was limited.

The research on the effects of practice on the transfer of prior knowledge suggests that practice facilitates the execution of the practiced task, but generalization to other tasks is limited. In fact, some researchers have posited that the amount of transfer to other tasks decreases as the amount of practice increases (VanLehn, 1996). Other work has indicated that practice can actually impair performance on tasks that resemble the practiced task but differ in how they are executed (Woltz, Gardner & Bell, 2000). Thus, the bulk of the research seems to support the specialized rather than the generalized view of practice. However, the effect of practice on transfer is not unequivocal. For instance, research in the area of device operation indicates that device operators can develop more flexible strategies for carrying out faultfinding tasks with increases in practice (Bibby & Payne, 1996). Furthermore, experts are less affected by superficial differences between problems than novices (Chi et al, 1981; Novick, 1988). While practice is not a sufficient condition for the development of expertise, it is a necessary component (Ericsson, 2001). One might suspect that the increased ability of experts to generalize across problems within their domain is partly due to their increased amount of practice within a domain.

However, the view that experts have the ability to generalize across problems within their domain has not gone unchallenged. While research has shown that experts are better able to generalize across problems within their domain than novices (Chi, 1981; Novick, 1988), other research has shown that even experts are capable of displaying limited transfer to problems in their own domain (Novick, 1988; Blessing & Ross, 1996). For instance, experts are affected by the superficial structure of the problem when the problem is easy to apply and execute (Novick, 1988), and experts are better at classifying problems that are typical for their domain than they are at classifying problems that are atypical (Blessing and Ross, 1996). Even more striking is the finding that experts often show poor transfer to tasks that are seemingly related to their level of expertise. For example, the exceptional memory of experts is specifically limited to their particular domain of expertise (e.g., Chase and Ericsson, 1982; Ericsson & Polson, 1988). In one study, a participant was able to recall 81 randomly presented digits within 200 hours of practice. This is almost 12 times the amount of digits that the average person can recall. However, when the participant was asked to remember a list of letters, his memory was limited to around seven letters (Chase and Ericsson, 1982)

In sum, the bulk research on practice and transfer indicates that the benefits of practice should only transfer to tasks that are very similar to the practiced tasks; transfer to superficially different tasks should be limited. However, there is some evidence that the effects of practice might lead to generalizations. This work is evident in the domains of device operation (Bibby & Payne, 1996) and expertise (Chi et al. 1981; Novick, 1988).

The missing link between practice and transfer might be found in the type of training that the user is given. For example, the review of the literature on learning from multiple examples suggests that learners can form generalizations when they are exposed to multiple examples of the principle. In effect, the provision of multiple examples is a form of practice. If one defines practice as two or more presentations or executions of the exact version of a procedure or task, then the presentation or execution of multiple examples (i.e., two or more different examples of the same principle) is merely practice with the exception that the practiced material occurs in different forms of the same principle, rather than the repeated presentation of an identical form of a task. This form of practice might lead to a more flexible understanding of a domain and consequently lead learners to become less affected by the superficial differences between problems. A similar argument has been proposed for explaining the difference between experts and novices (Ericsson, Krampe & Tesch-Romer, 1993). Suggesting that the type of practice is the essential factor for promoting transfer might solve the specialized verses general conundrum. Practice with a single task merely enables the learner to optimally execute the task (see Anderson, 1987). However, the effects of practice on transfer may be general when the type of training is varied (i.e., the provision of multiple examples).

The Transfer of Device Knowledge

My work over the last few years has been primarily concerned with the factors that can help promote the acquisition of novel device procedures. As mentioned in the preamble, the examination of transfer in the domain of device operation has several benefits. With a wide range of electronic devices appearing on the market in conjunction with the vast number of people using these machines, the study of device operation affords a great opportunity to study transfer as it occurs in the real world. Through the discovery of the underlying processes that govern device transfer, researchers can develop interventions to help people who are having difficulty with using novel devices. In this way, everyone can benefit from the use of technology. In the next section, I briefly outline the device-learning framework developed by O'Reilly and Dixon (2001a,b). Then, the research on the transfer of the procedural (O'Reilly & Dixon, 2001a) and the superficial aspects (O'Reilly & Dixon, 2001b) of a device is reviewed.

Device-learning Framework of O'Reilly and Dixon (2001a,b)

O'Reilly and Dixon (2001a,b) proposed a model for identifying and separating the factors that influence the transfer of device knowledge. In this framework, two forms of experience mediate transfer: specific and nonspecific experience. In many cases, a learner can immediately perform some of the steps for the new operating procedure based on their existing experience with similar devices. For example, analog tape players have controls that are often labeled *Play*, *Rewind*, and *Record*. With this knowledge at hand, a learner can directly use this control information to operate other tape players provided that they also know when to use the controls. The immediate application of such specific experience is referred to as *direct* transfer because an identifiable experience can be directly applied to aid in the acquisition of a related task without any additional learning. Specific experience occurs when features of the cue have been previously associated with features of the correct response; it implies that the user starts with a certain amount of information about the nature of the correct response in the current situation. In other cases, direct transfer could occur when there is an analogous relationship between two devices. For example, a compact disk player may have controls identified with icons like those found on the rewind and fast-forward controls on a tape player. One may be able to understand how to use such controls by forming an analogy between the controls on the tape player to the controls on the CD player. Direct transfer would occur in this situation as long as the order of steps required to execute the procedures were same for both the tape and the CD player. In other words, the episodes provide specific experience with the device, and the result is referred to as direct transfer.

In other circumstances, prior knowledge can aid in the acquisition of a novel procedure even when the experience is not directly relevant to the transfer task. For example, prior experience programming a VCR may be helpful in learning to program a different model VCR even if the two procedures share very few steps. The background knowledge of VCR's might provide the user with a familiar structure to organize and associate the steps to be learned. In turn, this structure could provide the user with a mnemonic for learning the new device more rapidly. For instance, most VCR's have controls for turning the power on, playing the tape, and adjusting the picture quality (tracking). Furthermore, one would expect to find these controls in specific locations. Often, the play, fast forward and rewind, and stop buttons are placed near one another, while the power and tracking controls are positioned in a different part of the interface. Such a familiar arrangement might help learners to encode the correct steps more easily and distinctively. The use of this experience might not influence performance at the outset, but the familiar locations of the controls should speed up the rate in which the task is learned. In this case, the learner benefits from nonspecific experience because there is no identifiable experience that can be directly used to perform the novel procedure. However, the prior experience can be utilized to make the task easier to learn. Nonspecific experience may cause the possible responses to be generally familiar, even if the user has little basis for selecting one over another. We assumed that such experience allows the responses to be represented more distinctively, so that there is less similarity between the representation of the correct response and other, incorrect responses. In this way, nonspecific experience may improve the rate of acquisition rather than immediate performance. The benefit of such nonspecific experience is referred to as *indirect* transfer because there is an improvement in the rate of task acquisition rather than the immediate application of specific knowledge. In other words, additional learning is required to perform the task without error, but transfer can be improved by utilizing the familiarity of the task.

It is likely that both direct and indirect transfer are involved in learning a new procedure. However, in our analysis the nature of their effects are distinct (O'Reilly & Dixon, 2001a,b). Direct transfer, in the limit, can make the acquisition of a procedure trivial because the required steps can be simply imported from a familiar device. On the other hand, indirect transfer merely provides the foundation for learning a procedure by providing a structure to acquire the procedure more quickly. In general, some further learning would be needed before the procedure could be carried out without error. Thus, we argued that when confronted with a new device, direct transfer reduces the likelihood of an error at the outset (i.e., initial error rate), while indirect transfer allows one to learn the procedure more quickly from the errors that do occur (i.e., learning rate).

We developed a tool for distinguishing whether a particular aspect of a procedure affected direct or indirect transfer (O'Reilly & Dixon, 2001a,b). In our model, the number of errors on a particular task was expressed by the following equation:

$$N = \frac{1-r}{s} \tag{1}$$

where r is the probability of performing a step correctly initially and s is the probability of learning a step after each error (e.g., Bower, 1961). The learning model is an extension of the all-or-none (one element model) described by Atkinson, Bower, and Crothers (1965) adapted to the operation of complex devices. We assumed that each step of the procedure is learned independently, that each step is either in a learned state or an unlearned state, and that on each trial, there is a fixed probability s of moving from the unlearned to the learned state. The equation can also be shown to be approximately true when distributed representations are used for the steps and connectionist learning algorithms are applied (O'Reilly & Dixon, 2001b).

If performance on each task in the transfer paradigm is measured as the number of errors, N, before reaching criterion, then transfer can be indexed by the improvement on the second task relative to the first, that is, N_1/N_2. Consequently, using Equation 1, the relative error rate, is

$$T = \frac{N_1}{N_2} = \frac{\frac{1-r_1}{s_1}}{\frac{1-r_2}{s_2}} \tag{2}$$

This expression can be transformed and simplified by taking the logarithm and rearranging the terms:

$$\begin{aligned} \ln T &= \ln(1-r_1) - \ln(s_1) - \ln(1-r_2) + \ln(s_2) \\ &= \left[\ln(1-r_1) - \ln(1-r_2)\right] + \left[\ln(s_2) - \ln(s_1)\right] \\ &= R + S \end{aligned} \tag{3}$$

where R is related to the change in the initial error rate from training to transfer and S depends on the change in learning rate. Thus, according to this analysis, effects on the initial error rate and effects on the learning rate should be additive when transfer is expressed as log relative errors. Using this framework, we were able to assess the aspects of the device that affected direct and indirect transfer using a logic similar to the additive factors technique (Sternberg, 1969). For example, if a manipulation interacts with a variable that is known to have an effect on the learning rate, then that manipulation is also likely to have an effect on the learning rate. In contrast, if a manipulation is additive with a variable that is known to an effect on the learning rate, then it is plausible to suppose that manipulation has its effect on the initial error rate. In other words, an interaction means that the two variables in question are apt to have their effects on the same transfer component. Alternatively, additive effects suggests that the variables in question have their effects on the distinct transfer components. Herein this procedure is referred to as the *selective effects* method.

Using this method, we were able to conclude that direct transfer was affected by the device's step order and control names; in contrast, indirect transfer was affected primarily by the device's interface and by the nature of the subprocedure labels (O'Reilly & Dixon, 2001a,b). More specifically, these conclusions were drawn from experiments that show the manipulations of subprocedure labels and step order are additive within one another, while the manipulations of interface and controls names were also additive with each other. Below, I describe the experiments that provide evidence for these selective influences on indirect and direct transfer. The experiments described in this section follow a similar format. Participants learned to execute a procedure by using a particular device, and then they attempted to operate another device that was related to the first device in some fashion. Transfer was assessed by comparing the number of errors made on the training task to the number of errors made on the transfer task. This type of design is referred to as a training and transfer paradigm.

In a training and transfer paradigm, we assessed the influence that the step order, subprocedure labels, and a hint had on a learners ability to transfer their prior device experience. Step order is simply the sequence of steps that are used to execute a task (e.g., first press *Start*, then *Transfer*, then *Hold*). Subprocedure labels refer to a written description of overall goal and function of the device procedure (e.g., Data Location Procedure). A hint is simply a statement that clues the user into the related structure of the training and the transfer tasks. We manipulated whether the step order and the subprocedure labels were the same or different from training to transfer. Participants were also either informed that the devices were similar or they were not given such a hint. We reasoned that the step order was the critical indicator for the process of direct transfer. If the steps of a procedure were rearranged from the training device then the learner should have great difficultly in directly applying their prior experience. Using the step order as a marker variable for direct transfer, we were able to assess the contributions of both the subprocedure labels and the provision of a hint on transfer. According to the selective effects method described above, if a variable interacts with step order, then it should have its effect on direct transfer; similarly, if a variable is additive with step order then it is likely that the variable has its influence on the learning rate and indirect transfer.

When participants were given a hint, transfer was greater when the steps remained intact and when the subprocedure labels were the same from the training to the transfer task. However, the effects of step order and the subprocedures labels were additive. In other words, subprocedure labels had an effect on indirect transfer and the learning rate. In contrast, the effect of step order was attributed to having an effect on direct transfer. Thus, the user's knowledge of the step order allowed them to reduce their initial error rate on the transfer task by providing them with specific experience that could be directly used to carry out the transfer task. On the other hand, the subprocedure labels allowed the user to learn the procedure more quickly because the subprocedure labels provided a structure for making the transfer task easier to learn. Although this experience could not provide the necessary information to carry out the procedure without error, the label did increase the learning rate by allowing the user to represent the correct steps and their orders more distinctively.

In contrast, when the participants were not provided with a hint, step order had almost no effect on transfer. In other words the hint manipulation and step order interacted. According to the selective effects method, the hint manipulation had its influence on direct transfer and the initial error rate. That is, when learners were instructed to use their prior experience, they were immediately able to apply their prior experience with the transfer device. In contrast, when the learners were not provided with such a cue, they failed to notice the relationship between the devices and consequently, they had to rely on other means for operating the device.

The poor transfer of the step order information was surprising because the devices had similar purposes and functions. With this additional level of similarity, one might expect the participants would notice that the devices were similar, and consequently, they should use this knowledge to help them execute the transfer task. However the use of the step order information only occurred when the participants were explicitly told that the devices were similar. We concluded that it was necessary for participants to receive explicit clues in order to notice the relational structure between novel devices. Without this reminding, the transfer of step order information was unlikely to occur. Thus, many learners do not automatically use their prior experience even when it is directly relevant to the task at hand. However, without a hint, there was still a benefit of providing the same subprocedure labels, suggesting that the labels affected an independent process of indirect transfer.

In sum, the findings from O'Reilly and Dixon (2001a) suggest that the transfer of prior device knowledge is influenced by the relationship between the structures of the devices (i.e., step order). That is, transfer is greater when the devices share a similar step order. Moreover, when the devices are novel, the transfer of the step order is dependent upon whether the participants are explicitly reminded of the similar operating procedure between the devices. However, even without a hint, prior knowledge can still promote transfer if the tasks share a similar subprocedure label. Finally, both the step order and hint affected the direct transfer and the initial error rate, while subprocedure labels had an effect on the learning rate and indirect transfer. Thus, both the hint and the step order allowed the user to reduce their initial error rate by directly applying their relevant prior experience to the transfer task. In contrast, the subprocedure labels helped the user to learn the procedure faster by providing a structure to assimilate the transfer steps and their respective orders.

Research on How the Physical Features of a Device Influence Transfer

In this line of research, we were interested in determining why the participants in the previous experiments failed to transfer their prior experience when they were not provided with a hint. We hypothesized that the dissimilar appearances of the devices used in O'Reilly and Dixon (2001a) might be the reason why spontaneous transfer was infrequent. When the learners are provided with a device that had a dissimilar appearance, they might fail to notice the relationship between the devices and, consequently, direct transfer is unlikely. However, when the appearances are similar the learners might spontaneously notice the connection, and as a result, they would be able to use the specific training experience to directly transfer their prior experience. The experiments by O'Reilly and Dixon (2001b) were designed to address the issue of whether the physical aspects of the device had any bearing on learning a new version of a device. More specifically, we examined whether the device appearance and interface layout had an affect on a participant's ability to learn a related device. The device appearance was defined the as the set of features that constitute the "look" of the device. These attributes included such things as the control name, gauges, indicators, fonts, and graphics that appeared on the machine. The interface layout was defined as the relative positions of the controls, gauges, and graphics on the interface.

The transfer of prior experience was investigated by training the participants to execute four tasks on four different versions of a simulated weather device. In the transfer phase of the experiment, participants were presented with the exact same four tasks as before, but the appearance of the devices were either the same or different from the appearance of the corresponding training devices. In this study, the participants were not provided with a hint that linked the related structure of the machines. Under these conditions, we hypothesized that the similar appearance would remind users that the training and the transfer devices were related. Once they noticed this relation, they could use their prior knowledge to operate the transfer device. In other words, the device's appearance should promote direct transfer and reduce the initial error rate. In contrast, when the appearance is different, the learners may not use their prior experience and direct transfer should not occur. As a second manipulation, we varied the interface layouts between the training and transfer devices. In one condition, the buttons in the transfer phase were positioned in the exact location as the controls used on the training device. In the different interface condition, the buttons were rearranged in the transfer phase. Despite the reorganization of controls, the actual order of the steps to achieve the correct task sequence was identical. We hypothesized that the interface layout would affect how quickly the learners would acquire the new device. Having a similar layout should enable the user to encode the steps more easily and distinctively. In other words, interface layout should affect indirect transfer.

The findings indicated that when the learners were not provided with a hint, the interface layout and the appearance had interactive effects on transfer. When the appearance was the same, a similar interface facilitated learning. However, when the appearance was different in the transfer phase, the interface layout had little effect on learning. We concluded that the device appearance reminded learners that the current machine was similar to a previously encountered device. Once this experience was reactivated, the participant was able to use their prior knowledge of the training task to transfer their experience to the novel device task. However, when the appearance was different, participants were unaware that their prior experience was related to the new task, and consequently they failed to use this knowledge to help them learn the novel device.

In a second study, we replicated the same experiment, but we provided the user with a hint that the devices were related. In this case, we expected that all of the users would be able to notice the correspondences between the devices and as a result, they should be able to directly transfer their prior experience. In other words, the interface layout should have an effect even when the appearance was different because the participants would be aware that their prior experience was relevant to the transfer task. The findings supported our hypothesis: The appearance and the interface layout had additive effects on transfer. The hint allowed the participants to immediately apply their prior experience even when the devices had a dissimilar appearance; without this clue, the learners could not use their relevant experience of the interface layouts. Thus, the hint affected direct transfer. However, the hint also affected indirect transfer because the hint manipulation interacted with interface layout. Thus, the hint allowed the user to directly apply their prior experience (direct transfer) and it also helped the user to learn the task faster by promoting indirect transfer.

In sum, we attributed the effect of appearance to direct transfer, while the effects of interface were attributed to indirect transfer. The provision of a similar appearance reminded the user that the machines were related, and consequently, the participant could directly use this information to reduce their initial error rate on the transfer task. However when the user was given an explicit hint, they were able to both directly and indirectly transfer their prior experience irrespective of the device appearance. In contrast, appearance was additive with interface layout, suggesting that the interface layout affected the process of indirect transfer. Having the controls in a similar configuration made the controls and the steps easier more distinctive, and consequently, the steps were easier to learn. Thus, the influence of interface layout was ascribed to indirect transfer and an increase in the learning rate.

If the participants had acquired the deep elements of the task, then the reorganization of the controls and the different appearances should have had no bearing on their ability to learn a new device because the order of the steps within a procedure remain intact. The position of the controls and the appearances are merely superficial aspects of the task, and changing these superficial elements should not influence performance because the essential aspects of the task are preserved (e.g., meaning of the controls, step order and task goals). Nevertheless, when participants received a similar interface, they made considerably fewer errors than when they were given a dissimilar interface layout. In Experiment 2 of O'Reilly and Dixon (2001b), this effect was independent of the effect of appearance: Appearance and interface had separate (additive) influences on the participant's ability to learn a new device. Thus, according to the model, interface layout affected the learning rate and the process of indirect transfer, while the appearance and hint affected the initial error rate and the process of direct transfer.

The results of Experiments 1 and 2 (O'Reilly & Dixon, 2001b) indicated that the superficial appearance of a device can influence a participant's ability to learn a novel device. However the findings of the first two experiments were confounded: The appearance manipulation contained both the graphical content (type of gauge, font, and graphics) and the type of control name. In the different appearance conditions, the control labels were synonymous in meaning. For instance, the controls on one device might be labeled *Insert, Calculate,* and *Distance* while the controls on the related device might be labeled *Enter, Compute*, and *Length.* While the words are similar in meaning, they are different in appearance. In the same appearance condition the names between the training and transfer devices were identical *Enter*, and *Enter*. It might have been the case that the identical control names were responsible for the effect of the appearance manipulation in Experiments 1 and 2 and not the graphical content per se.

In Experiment 3, the graphical content and the control names of the devices were manipulated separately. Participants were asked to operate transfer devices that had control names that were identical to the training control names or the names were synonymous. As a second manipulation, the graphics could be either the same or different from the graphics used on the training device. This design allowed one to determine the effect of the graphical content irrespective of the control names. In this experiment, the graphical content had no effect on a participant's ability to learn a new device. However, the control labels did have a large influence: When the control labels were identical from training to transfer, participants made fewer errors than when they were given synonymous control names. Consequently, we attributed the affect of the appearance manipulation in the previous studies to the control labels. When learners are not cued to use their prior experience, an identical control name can remind users that the devices are related. Once the user has been reminded, they can retrieve their prior experience of *when* a control should be executed in the context of the procedure. However, when the names are synonymous, the user does not notice that the synonyms are related to the controls on the training device. As a consequence, the learner does not use their prior experience (i.e., no direct transfer) and transfer is low. In this case, the prior experience is the knowledge of when to instantiate particular controls in terms of the overall procedure (step order). In short, direct transfer requires both the memory of the control name, and the memory for when to use that control in the procedure. When the names are synonymous, the user does not automatically connect the transfer control names with the control names stored in memory. When this connection fails, transfer gains do not occur because the prior experience of when to use the step is not accessed.

Finally, it is important not to confuse the four training procedures used in O'Reilly and Dixon (2001b) as multiple examples. In the problem solving literature, a multiple example is defined as two or more versions of a problem that illustrate the same principle (Gick & Holyoak, 1983). That is, multiple examples are based on one or more superficially different versions of a problem that contains the same structure. In the domain of device operation, the structure of the task corresponds to the procedural goal and the steps used to execute the task. Recall that each of the four training tasks in the above experiments involved a different goal and a different set of steps to carry out the goal. Consequently, the four training devices are not multiple examples. Thus, in the experiment of O'Reilly and Dixon (2001b), there is only one example of each of the structurally dissimilar procedures. In order for multiple examples to occur in this paradigm, the user would have to be provided with two or more superficially different devices that carried out the same task. This issue is taken up in Experiment 4 of this dissertation.

In sum, the research on the transfer of prior device knowledge indicates that novices are often influenced by the superficial structure of the device (O'Reilly & Dixon, 2001b). In turn, when the superficial features of the device are different from a previously encountered device, learners often require a hint in order for them to utilize their prior experiences (O'Reilly & Dixon, 2001a; O'Reilly & Dixon, 2001b, Experiment 2). In cases where the participants are not explicitly informed about the relationship between the devices, the control names can remind them of the relationship. However, this cueing is only evident when the control names are identical and not when they are synonymous (O'Reilly & Dixon, 2001b Experiment 3). This suggests that novice users process devices and their procedures at a very superficial level. If processing occurred at a deep level, one would expect the users to transfer their prior experience when the appearance is different and the control names are synonymous based on the similar meaning of the controls.

However, the experiments show that learners are able to utilize some of their prior experience even in the absence of a hint: When participants are provided with the same subprocedure labels as in training, transfer improves (O'Reilly & Dixon, 2001a). In this case, the user is able to learn the transfer task faster because the label provides a background for making the steps easier to learn. Finally, the step order, control names, affected the initial error rate and direct transfer, while the subprocedure labels and the interface layout affected the learning rate and indirect transfer. In O'Reilly and Dixon (2001b) the hint was found to influence both direct and indirect transfer. The user can reduce their initial error rate on the transfer task by simply importing their specific experience that is relevant in executing the task. Conversely, if the user has no specific experience that is relevant to the transfer task, they can increase their learning rate by integrating the steps and order information into the context of the subprocedure label, or the interface layouts.

Before the experiments in the next section are proposed, one criticism of the research on superficial similarity is addressed. Some researchers have argued that the effects of superficial similarity on the transfer of prior knowledge are simply due to a change in the context from the training to the transfer conditions of the experiments (Needham & Begg, 1991; Spencer & Weisberg 1986). The researchers point out that in many of the earlier studies of analogical transfer, participants were required to approach one task in the context of memorizing story, while in the transfer phase they were required to solve a problem. For instance, Needham and Begg (1991) found that transfer was greatly reduced when one of the training or transfer conditions contained a memorization task and when the other contained a problem-solving task. However, the experimenters found that transfer was much greater (the effect of superficial similarity was lower) when the training and transfer tasks were both with the same task goals (i.e., memory task in training/ memory task in transfer, or problem solving task in training/ problem solving task in transfer.). In our experiments, this argument does not apply because participants are required to solve both the training and the transfer tasks in the same modality (i.e., problem solving). Furthermore, the devices used in both the training and transfer tasks were from the same domain: weather devices. In other words, the effect of superficial similarity observed in O'Reilly and Dixon (2001a, b) was not be due to a change in context but rather a change in the superficial nature of the device itself.

Current Experiments and Predictions Based on the Review of the Literature.

The purpose of the present investigation is to determine the effect of varied training (multiple examples) and practice on the transfer of device experience. More specifically, I am interested in whether the combined manipulations of practice and the provision of multiple examples can help learners to spontaneously transfer their prior knowledge to a superficially dissimilar device. Without a hint, participants rarely transfer their experience when the devices look different (O'Reilly & Dixon, 2001b). To the best of my knowledge, no study has investigated both the effects of multiple examples and practice together. On the surface, one could argue that it is unnecessary to examine the effect of practice and varied training because the individual effects of these manipulations are already clear in the literature: Varying the type of training should reduce the effects of superficial similarity (see review on learning from multiple examples), and increasing the amount of practice should not reduce the effects of superficial similarity (see review on the effects of practice on transfer). The latter result is predicted because the benefits of practice are generally assumed to be limited to the practiced device, and transfer to superficially different devices should be minimal.

However, the combination of the two types of training might produce different effects than their individual contributions suggest. That is, the manipulations may interact when they are varied together. For instance, the specialized effect of practice on transfer might be reduced when an individual is also provided with multiple examples of a device. The provision of multiple examples may allow the user to develop a more flexible representation (i.e., less susceptible to superficial dissimilarity) of the machine, and this representation could be strengthened by practice. Furthermore, it is not entirely clear whether the effects of practice are completely specialized. For instance, in their analysis of the effects of practice on device learning, Bibby and Payne (1996) found that participants became more efficient in carrying out a faultfinding procedure with increased practice. After practice, there was a switch from using mechanical strategies (e.g., sequentially checking all possible broken components in a fault finding task) to a more sophisticated search-reducing strategies (e.g., restricting the possible broken components). This finding suggests that practice can enable a user to discover better strategies for executing a task (also see Lee & Anderson, 2001). It is therefore possible that practice could facilitate the transfer of prior device experience to superficially dissimilar devices.

Other work has suggested that learners can develop general rules with practice and that these rules allow the learner to generalize over examples (Anderson, Fincham &

Scott, 1997). Moreover, the research on expertise suggests that experts are able to generalize across superficially different problems within their domain (Chi et al. 1981, Novick, 1988). While expertise is not simply as function of increased practice, practice does play an important role in the development of expertise (Ericsson et al, 1993). With increasing practice, one might expect learners to develop some abstract knowledge concerning the operation of devices and that this knowledge could better enable them to generalize (i.e. transfer) across superficial instances of a novel device. In sum, it is unclear as to how both practice and the provision of varied examples will affect the transfer of prior knowledge.

In addition, it is also unclear whether practice will have its effect on indirect transfer and the learning rate or direct transfer and the initial error rate. On the one hand, one might assume that practice with a particular device provides specific experience that can be immediately applied to a transfer task. That is, the user should be able to operate the machine without any further learning, and as a result, the initial error rate will be reduced, and direct transfer will occur. For instance, one might expect that with practice the learners might be able to learn the names of the controls. When they encounter a novel device that has same names, they should be able to apply this learning on demand. On the other hand, one might suspect that practice could provide non-specific experience that would lead to a faster learning rate and indirect transfer. For instance, with increases in practice, the procedure and position of the controls on the training device might become more familiar. When the learner is exposed to a device with a similar procedure and configuration, they might be able to learn the new interface faster because the non-specific experience allows the correct steps to be more easily encoded and retained.

Further, it is also unclear how the provision of multiple examples will influence transfer. On the one hand, multiple examples might allow the user to increase the learning rate by allowing the user to become more familiar with the meaning of the control names. This familiarity should makes the control names easier to learn, but it will not provide error free performance because it is assumed that practice does not also correspond to an increase of knowing when to use a step during the procedure (step order). Alternatively, the provision of multiple examples might provide specific experience that enables the user to directly transfer their prior experience. That is, practice would increase the chances of acquiring the meaning of the names and when to instantiate the controls (step order).

In sum, it is unclear as to how both practice and the provision of varied examples will affect the transfer of prior knowledge, and whether practice and multiple examples will influence direct or indirect transfer. Thus, further experimentation is required to resolve the issue. In the next section, I describe four experiments designed to assess the effects of practice and multiple examples on the transfer of prior device experience. Experiments 1, 2 and 3 examine whether the practice influences direct transfer and the initial error rate, or indirect transfer and the learning rate. Finally, the fourth experiment assesses the effect of multiple examples and practice on the transfer of prior experience.

Chapter 2

Experiment 1

The purpose of this experiment was to determine whether practice has any effect on a participant's ability to transfer to a superficially different device. Previous work has indicted that device novices are influenced by the superficial nature of the control names and the interface layouts (O'Reilly and Dixon, 2001b). The goal of this experiment was to establish whether practice could increase the tendency for participants to transfer their knowledge when the names are synonymous and when the interface layouts are different. Furthermore, I wanted to examine whether practice affects indirect or direct transfer. Other work has indicated the control names have an influence on the initial error rate and direct transfer while the interface layouts have an effect on indirect transfer and the learning rate (O'Reilly and Dixon 2001b).

The devices used in this study were simulated on a computer screen and learners interacted with the device by using the computer mouse to "press" the device buttons. All of the experiments reported here had the same structure: participants learned the operating procedures for eight different devices, with each of the last four being related in some way to one of the first four. The first four devices are referred to as training devices and the last four as transfer devices. Performance on each of the four transfer devices was examined as a function of the relationship to the corresponding training device. This design is similar to the methods utilized by O'Reilly & Dixon (2001b).

The experiment was a 2 x 2 x 2 Mixed design with practice and interface layout as within-subjects variables and control names as the between factor. The interface layout was either the same or different and the control names were either identical or synonymous. In addition, and practice was either high or low. In the high-practice condition, participants were required to complete the procedure 10 times. In the low-practice condition, participants were required to complete the procedure twice. In the transfer task, all of the devices were operated for two trials. Participants first learned the four training devices and then they attempted to operate the four transfer devices.

The dependent measure was the log number of errors the participants make on the transfer task relative to that on the training task (i.e., Transfer = ln [training errors +1]/[transfer errors +1]). More specifically, the relative number of errors for each participant was calculated by dividing the number of errors on the first task by the number on the second; this helps to minimize the impact of individual differences. A remaining issue, though, is that the distribution of relative errors in such tasks are often skewed. To address this problem, the relative error scores were transformed by taking the logarithm. Taking the logarithm reduces the skew and is common in designs such as these (Bibby, & Payne, 1993, 1996; Schumacher, & Gentner, 1988). As well, using log relative errors as a measure of transfer facilitates interpretation of the effects because, as shown in Equations 2 and 3, effects on initial error rate and learning rates should be additive. Some participants made no errors on either the training or task or the transfer task, so 1 was added to each participant's training and transfer scores before computing the overall transfer score. This was necessary in order to avoid undefined overall transfer measure.

According to the selective effects method described above, if practice interacts with a variable known to affect direct transfer (i.e., control names) practice should also have an effect on direct transfer. In contrast, if the practice is additive with control names, then it is likely that practice has its primary influence on indirect transfer and the learning rate. In a similar vein, O'Reilly and Dixon (2001b) concluded that interface layout affect has an effect on indirect transfer and the learning rate. Thus, if practice interacts with interface layout, then it is likely that practice also has an effect on indirect transfer and the learning rate. In contrast, if practice is additive with interface layout then one might attribute the effects of practice to direct transfer and the initial error rate.

Furthermore, if the effects of practice were general, then one would expect to see more transfer in the synonymous-control-names / high-practice condition than the synonymous-control names / low-practice condition. In this case, the user should be able to generalize across the synonymous control names. Moreover, if the effects of practice were completely general, one would not expect to see a difference between the synonymous and the identical names condition for the high-practice conditions (see

Figure 4, pg. 67). In contrast, if the effects of practice were specialized, then one would expect to see an advantage for the identical names over the synonymous names condition even when practice is high. In this situation, the high-practice condition should be superior to the low-practice condition only when the names are identical because practice is assumed to be optimal (specialized) only for the practiced task. This logic also applies to the interface layouts: If practice is general one should not expect to see an effect for interface when practice is high; in contrast if the effects of practice are specialized there should be difference between the interface conditions when the practice is high.

Before the experiment is described in more detail, it is important to reiterate a point made earlier: The presentation of four training devices that carryout different procedures is not considered as multiple examples. Multiple examples are defined as the presentation and execution of 2 or more *versions* of the *same* procedure. For example, a learner is said to have executed "multiple examples" when they complete two or more versions (e.g., with different graphics, layouts etc.) of a device that carry out the same procedure (e.g., Startup Procedure). However, the procedure for each training device in the first two experiments reported here are different, and therefore, the training devices do not constitute multiple examples.

Accordingly, practice as is it is manipulated in the first two experiments, is also not considered as multiple examples. In these experiments, practice was defined as the repeated presentation and execution of a single procedure on the identical version of a device. For example, "practicing" a device would mean that a learner completes a given procedure (e.g., Startup Procedure) two or more times on the same device (i.e., with identical graphics, gauges, control names position of the buttons, etc.). Although all four training devices were weather machines, the procedures for each device were entirely different (e.g., Startup Procedure; Sensor Activation Procedure, Immediate Reading Procedure, and Data Transfer Procedure). Changing the overall goal and the steps used to achieve that goal effectively changes the principle or structure of the task and consequently violates the above definition of multiple examples. Thus, the repeated presentation of the same procedure (practice) is not considered as multiple examples. The issue of multiple examples is explored in Experiment 4.

Method

Participants

The participants were 48 introductory psychology students from the University of Alberta who were given course credit for their participation.

Materials

Thirty-three hypothetical weather-monitoring devices were simulated on a computer screen. Figure 1 is an example of a typical interface used in this research. The device interfaces consisted of a distinctive dialog box with labeled buttons; participants "operated" the devices by using a computer mouse to click on the simulated buttons in a

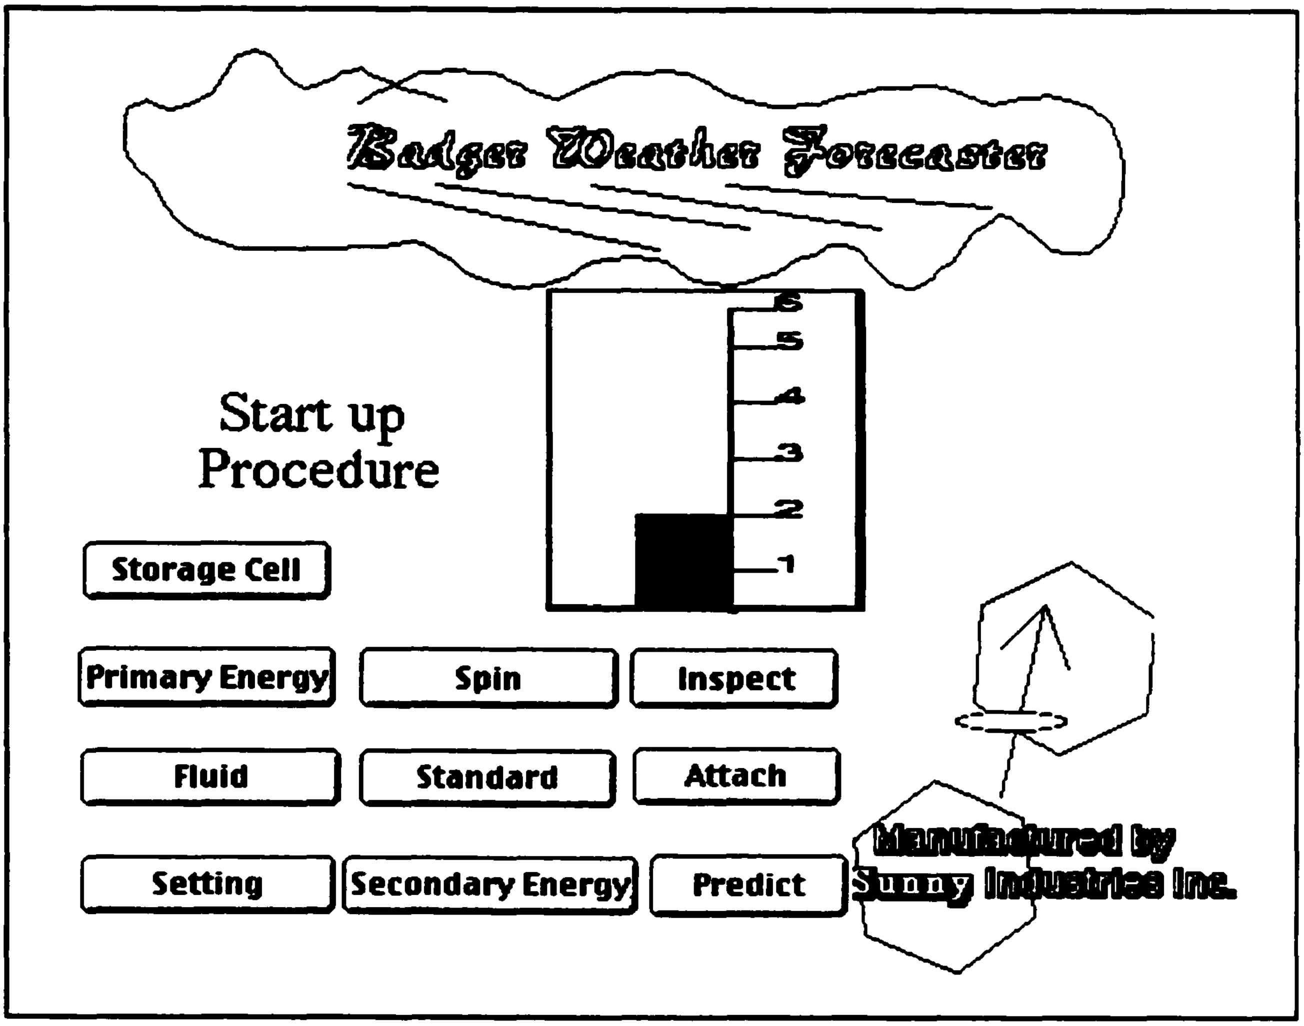

Figure 1. A typical device used in these experiments.

particular order. All device consoles contained a set of ten buttons, each of which was necessary to carry out the designated procedure. One device was used as a warm-up device, while the remaining 32 were different versions of 4 different weather devices each with its own unique operating procedure. The procedures were be labeled as "Start-up Procedure," "Sensor Activation," "Data Location," and "System Cleaning." Each procedure had a unique goal, controls, and sequence of steps used to carry out the task. The name of the operating procedure was indicated on the device console. As well, a brief description of the procedure appeared on the screen prior to the presentation of the device itself. This description contained an overview of the purpose of the procedure, but it did not contain any information on how to execute the task (i.e., the sequence of steps). For each of the four procedures, there were two synonymous control names (e.g., *temperature/ thermal*), and two interface layouts. The appearance of each device consisted of a graphical logo, a gauge used for operation feedback, a manufacturer name, and button labels. The appearance (i.e., gauge, graphics) of the devices for a given procedure was identical.

Participants were provided with a two-page instruction booklet that described the nature of the devices, the number of devices they were required to learn, and the consequences of making both correct and incorrect button sequences. The instructions also explained that each device was operated by pressing 10 buttons in a specific sequence, and the participants' task was to find the correct 10-button sequence for that particular procedure.

Design and Procedure

In the low-practice conditions, the participants were required to complete the task twice. In the high-practice conditions, participants were required to complete the task 10 times. In both the high and low-practice conditions, the successful completion of the procedure meant the learner progressed through the complete procedure irrespective of the number of errors that they committed by the time they reached the last step. That is, the participants had to perform the steps in the correct order before they progressed onto the next trial, regardless of whether they made any errors on that procedure. First, participants learned a warm-up task using a device interface that was comparable to the experimental devices to familiarize them with using the mouse and interacting with the device. This was followed by four training and four transfer tasks. The transfer devices had the same labeled procedures in the same order as during training. However, the control names were either identical or synonymous, and the interface layout was either the same or different. In the transfer task, all devices were operated for two trials.

The data in this experiment were originally collected as two separate experiments: In one experiment the participants were provided with transfer devices which had names that were identical to the corresponding training names, and, in the other experiment, the corresponding names were synonymous. The synonymous name experiment was run first, followed later by the identical name experiment. The data are analyzed and reported together for clarity of exposition and inference. Despite the potential problems with using this design, (e.g., lack of random assignment to the names conditions) this approach should not compromise the interpretation of the results. First, the conditions were run within a relatively short time of each other, and, second, the participants in both conditions came from the same population of introductory psychology students for the academic term. All four combinations of interface and amount of practice were used for each participant (i.e., the within-subjects factors). The order of the four transfer conditions (i.e., interface and practice) was balanced across subjects, and the particular devices, procedure, and versions were randomly selected for each subject and condition.

Before beginning each task, all participants were asked to carefully read a two-page instruction booklet. After reading the instructional materials, participants were encouraged to operate the device. When a participant made an error, feedback was given regarding the correct alternative, and the task was continued from that step. For instance, suppose that the user was required to press a series of buttons in the following order-*Guard Mode*, *Infrared*, *Movement*, *Test*, and *Activate*. If the participant pressed the button *Delay*, the computer provided them with a note that said, "You should have pressed the *Guard Mode* button". At that point, the learner had to press the button indicated in the message (*Guard Mode).* For example, if the user continued by correctly pressed *Guard Mode*, then *Infrared*, then *Movement*, but then an incorrect control

Activate, the computer would indicate that he or she should have pressed the *Test* button. After the user pressed the *Test* button, he or she was required to continue with the remainder of the procedure; in this case, he or she would be required to press *Activate*. A single trial was completed when the participant reached the last step in the procedure. An interface gauge increased in value after each correct button press.

Results

The average number of errors for the training and transfer tasks in each condition are shown in Table 1. The scores were tabulated by summing the number of errors each participant made on the first two complete trials of both the training and transfer tasks. This measure was taken to ensure an unbiased comparison between the low and high-practice conditions.

Table 1

Raw errors and standard errors for the training and transfer tasks as a function of interface layout, control names and practice

Synonymous Names

	Training Task				Transfer Task			
	Same Layout		Different Layout		Same Layout		Different Layout	
High-practice 10	8.7	(0.53)	7.8	(0.55)	3.2	(0.64)	7.4	(0.60)
Low-practice 2	8.1	(0.54)	8.5	(0.40)	6.3	(0.57)	8.3	(0.58)

Identical Names

	Training Task				Transfer Task			
	Same Layout		Different Layout		Same Layout		Different Layout	
High-practice 10	9.7	(0.52)	8.6	(0.78)	1.2	(0.40)	5.3	(1.05)
Low-practice 2	9.1	(0.64)	8.9	(0.64)	4.6	(0.71)	6.5	(0.86)

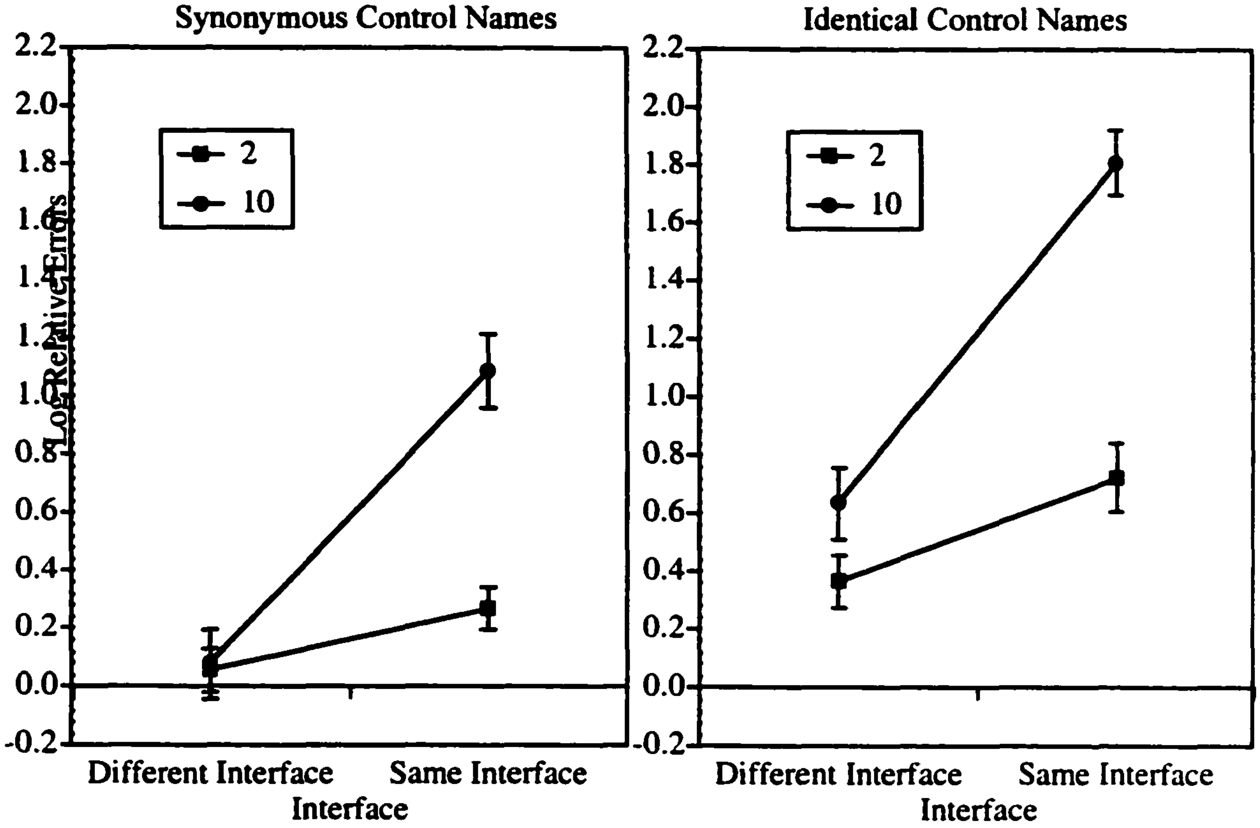

Figure 2. Transfer as a function of practice, interface layout and control names.

Figure 2 shows the amount of transfer (i.e., log relative errors) as a function of practice, interface, and control names. Higher scores indicate more transfer. The results show that control names, interface layout, and practice all influenced the amount of transfer. In particular, providing the same interface layout or the identical names improves the transfer of training. However, practice seems to have a specific influence only on learning the interface layouts and, to a lesser extent, the control names. This is evidenced by the large apparent interaction between control names and interface layout, and the seemingly small interaction of practice and control names. Below, these impressions are examined by comparing the relative fits of various models using the maximum likelihood ratio.

The traditional approach for evaluating experimental data is to perform null hypothesis significance testing. For example, the claim that the effect of control names is additive with the interface layout manipulation would be evaluated by testing the interaction term in an analysis of variance. Finding a reliable interaction would entail rejecting the null hypothesis that the interaction was zero and the effects were additive. Despite its common usage, this approach has a number of well-known problems (e.g., Cohen, 1994; Loftus, 1993). One category of problems is that, on its own, rejecting or failing to reject the null hypothesis does not provide definitive evidence that a theoretically important effect does or does not exist. On one hand, a null hypothesis can be rejected on the basis of effects of a trivial magnitude when the power is very high (e.g., Thompson, 1993); on the other hand, when power is weak to moderate, one can easily fail to reject the null hypothesis even when the magnitude of the effect is substantial. An alternative approach suggested by Dixon and O'Reilly (1999) is to identify the different possible interpretations of the results that would be theoretically meaningful and then assess how well those interpretations match the data.

Below, the evidence for a model is examined by comparing its fit to several plausible alternatives. An argument is provided for why each alternative might be theoretically defensible and the maximum likelihood ratio is calculated, λ, as a concise summary of the relative fits of the models. The likelihood ratio indicates how likely the data are given the best fit of one model relative to how likely the data are given the best fit of the alternative. The maximum likelihood ratio is closely related to the sums of squares obtained with linear modeling techniques (such as analysis of variance). The likelihood of n independent normal observations, given a particular model, is:

$$\prod \frac{1}{\sqrt{2\pi\hat{\sigma}^2}} \exp\left(\frac{(X_i-\hat{\mu}_i)^2}{\hat{\sigma}^2}\right) = \left(\frac{1}{2\pi\hat{\sigma}^2}\right)^{\frac{n}{2}} \exp\left(\frac{\sum(X_i-\hat{\mu}_i)^2}{\hat{\sigma}^2}\right)$$

where $\hat{\sigma}^2$ and $\hat{\mu}_i$ are the maximum likelihood estimates of the variance and the means under the model. The maximum likelihood estimate of the variance is

$$\frac{\sum(X_i-\hat{\mu}_i)^2}{n} = \frac{SSE}{n}.$$

However in a mixed design, the observations for each subject are dependent. Thus, the likelihood can be rewritten as product of likelihood of n independent subject means with degrees of freedom n/2 and a likelihood of n(c-1) independent degrees of freedom for the within subjects variance:

$$\left(\frac{1}{2\pi}\right)^{\frac{n(c-1)}{2}}\left(\frac{1}{SSE_w}\right)^{\frac{n(c-1)}{2}}e^{n}\left(\frac{1}{2\pi}\right)^{\frac{n}{2}}\left(\frac{1}{SSE_b}\right)^{\frac{n}{2}}e^{n}$$

Thus, the likelihood of two such ratios for two different models can be written as:

$$\lambda=\left(\frac{SSE_{2w}}{SSE_{1w}}\right)^{\frac{n(c-1)}{2}}\left(\frac{SSE_{2b}}{SSE_{1b}}\right)^{\frac{n}{2}} \tag{3}$$

where SSE_1 and SSE_2 reflect the residual error sum of squares in the fits of the two models, n is the number of participants and c is the number of conditions (Dixon & O'Reilly, 1999). The subscript w reflects the contribution of the within-subjects variance, while the subscript b reflects the contribution of the between subjects variance. In many situations, evaluating the likelihood ratio is similar to testing null hypotheses, and in some prototypical hypothesis testing situations, obtaining a likelihood ratio of 10 or greater corresponds approximately to rejecting the null hypothesis with $\alpha<.05$ (Dixon, 1998).

Effect of Practice. One possibility is that there is an effect for interface layout and control names but no effect for practice. This model is motivated by the finding that both the control names and the interface layouts influence the transfer of device knowledge (O'Reilly & Dixon, 2001). An effect for practice is not predicted because it is assumed that once a participant learns a procedure, further practice on a device will not improve performance when the user is asked to operate the identical device in the transfer task. In a similar vein, increases in practice on a single device should not influence transfer to superficially related devices because it is assumed that the effects of practice does not generalize to superficially dissimilar devices. However, when a model that contains effects for interface and control names was compared to a model that predicts an effect for all three factors, λ was over 1000, indicating strong support for a model that includes an effect for practice above the effects of interface and control names.

Effect for Control Names. Another possibility is that there is an effect for practice and interface layout, but no effect for control names. An effect of practice is predicted because practice has been shown to effect performance in a wide variety of domains (e.g., Crossman, 1959; Kolers, 1976). Furthermore, the above comparison indicates that practice does have an influence on the transfer of device experience. An effect for interface layout is predicted because it is assumed that the users will operate the device by remembering the positions of the controls rather than by remembering the name of the controls. This prediction is based on the assumption that the control names on the device are uninformative of the device's function, and therefore, remembering the sequence in terms of the relative locations of the controls is the best way to learn the procedure. However, when a model that contains effects for interface and practice was compared to a model that predicts an effect for all three factors, λ was over 1000, indicating strong support for a model that includes an effect for control names in addition to the effects of interface and control names.

Effect for Interface Layout. A third possibility is that there is an effect for practice and control names, but no effect of interface layout. An effect of practice is predicted because prior research indicates that practice has effects in many domains (e.g., Crossman, 1959; Kolers, 1976). An effect of names is predicted because it is assumed that the user operates the devices by remembering the names on the controls, rather than by remembering the layouts of the controls. An effect of interface is not predicted because the layout of the controls is not considered relevant to the procedural goal. Consequently, it is assumed that the user executes the task on the basis of the semantic information contained in the controls. When a model that contains effects of practice and controls names is compared to a model that contains an effect for all three factors, λ was over 1000, indicating strong support for a model that includes an effect for interface layout names above and over the effects of interface and control names.

Interaction between Interface Layout and Practice. One motivation for the current experiment was to determine whether practice had an influence on direct or indirect transfer. In accordance with the selective effect method, if practice interacts with a variable known to affect indirect transfer (i.e., interface layout), then it is likely that practice also affects indirect transfer. To test this idea, a model that predicts additive effects for practice, names and interface layout was compared to a model that also predicts these effects, but in addition, includes an interaction between the layout and practice. When the simpler additive model was compared to the model that also includes an interaction, λ was over 1000, indicating that the data provide strong support for a model that includes an interaction between interface layout and practice over a model that does not include such an interaction. Thus, it is likely that practice affects the process of indirect transfer and the learning rate.

Interaction between Control Names and Practice. In line with the above comparison, if practice is additive with the control names (a variable that seems to influence direct transfer), then it is likely that practice has an effect on indirect transfer. Consequently, a comparison was made between a model that included effects for control names, interface and practice to a model that included these effects plus an interaction between control names and practice. When the simpler additive model was compared to the model that also includes an interaction, λ was 4.42, indicating that the data provide little support for a model that includes an interaction between control names and practice relative to a model that does not include such an interaction. Positive evidence for additive effects (rather than simply the lack of evidence for an interaction) can be obtained by comparing the additive model to a model that compares an interaction of a particular size. In this experiment, an interaction with a magnitude of 20% (of log relative errors) was considered to be theoretically important; this approximately translates into an interaction of 0.2 in log relative errors. By way of comparison, the size of the interaction between the interface and control names was 0.20. When the additive model was compared was compared to a model that predicts an interaction with a magnitude of .20, the value of λ was 2.84. This indicates that the data provide weak evidence for the additive model over the interactive model.

Discussion

The results of Experiment 1 provide several important findings. First, the superficial manipulations of control names and interface layout both have an influence on transfer: Transfer was improved by either providing the identical control name or by providing a similar interface. This finding replicates the results of O'Reilly and Dixon (2001) who also provided evidence for an influence of superficial similarity on a learner's ability to transfer their prior experiences. Second, practice depends on the interface layout: Practice greatly improves transfer when the interface on the transfer device is the same, while practice only has little effect when the interface is different. This result supports the specialized view of practice which asserts that practice helps a participant execute the practiced task more efficiently but that this benefit does not readily transfer to a related task with a different interface layout.

More importantly, the data suggest that practice influences indirect transfer. According to the selective effects method, any variable that interacts with a variable known to affect a given component of transfer the variable in question is also likely to have an effect on the same component. In this experiment, practice interacted with the interface layout suggesting that practice also influences indirect transfer. Thus, it is likely that practice helps improve performance by increasing the rate in which the transfer task is learned. A familiar configuration of controls could provide a structure and organization to learn the steps and their orders more easily. Such a configuration would improve the rate of acquisition by making the correct responses familiar and more distinctive. With increases in practice, transfer is improved because the association between the interface layout and correct responses become even more familiar and distinctive.

Although the interaction between practice and interface suggests that practice has an influence on indirect transfer, the data for the interaction of practice with the manipulation of control names is unclear. If practice affected indirect transfer then one would also expect to find an additive pattern of results between practice and names. This is because prior work has suggested that the control names influence direct transfer.

According to the selective effects method, if practice affects indirect transfer then it should be additive with control names. However the likelihood ratios provided only weak evidence for an additive model over a model that includes an interaction with a magnitude of .20. One reason for this could be due to the manner in which the control names were manipulated. Recall that the controls names in this experiment were manipulated between subjects. The failure to find definitive evidence for either an interaction or additive effects might simply be due to the increased variance associated with the between subjects design. If this were the case, then one might expect to find clearer evidence if the control names were manipulated within subjects.

Chapter 3

Experiment 2

The purpose of this experiment was to determine whether the effect of practice in the previous experiment could be attributed to indirect transfer. This experiment was a replication of Experiment 1 with the exception that the control names were manipulated within subjects, while the interface layout variable was manipulated between subjects. The clear interaction of practice and interface layout indicated that practice has its main influence on indirect transfer. However, the pattern of data for the control names and practice was less clear: There was little evidence for an additive pattern between names and practice. In turn, there was also no reason to expect that the pattern was interactive. This experiment was designed to resolve this issue by manipulating the control names variable in a more precise within subjects design. If the pattern of results in this experiment are additive with (between practice and control names) then one can conclude that the effect of practice is likely due to indirect transfer. Furthermore, if the overall pattern of results in this experiment matches up with the overall pattern obtained in Experiment 1, then such a replication would lend more support to the conclusion drawn from the data. In particular, similar pattern of data would provide converging evidence that practice influences indirect transfer and that the effect of practice are specialized.

Method

Participants

The participants were 48 introductory psychology students from the University of Alberta who were given course credit for their participation.

Materials

The materials were exactly the same as those used in the previous experiment. Thirty-three hypothetical weather-monitoring devices were simulated on a computer screen. One device was used as warm-up device, while the remaining 32 were different versions of 4 different weather devices each with its own unique operating procedure. Participants were also provided with the same instructional booklet as in Experiment 1.

Design and Procedure

The design and procedure in this experiment was the same as the procedure used Experiment 1 with the exception that the names and the practice factors were manipulated within subjects, while the interface layout variable was manipulated as between subjects. Thus, for half of the procedures, participants were given high practice, and for the other half they were given low practice. Similarly, for half of the transfer devices, the participants were provided with control names that were synonymous to the training names, while the other half of the devices contained names that were identical to the training names.

The data in this study were originally collected as two separate experiments: In one experiment, the participants were provided with transfer devices that had interface layouts that were same as the corresponding training interface layout, and, in the other experiment, the corresponding interface layouts were different. The same-interface experiment was run first, followed later by the different-interface experiment. The data are analyzed and reported together for clarity of exposition and inference. This approach should not compromise the results because the conditions were run within a relatively short time of each other and, the participants in both conditions came from the same population of introductory psychology students for the same academic term. All four combinations of control names and amount of practice were used for each participant (i.e., the within-subjects factors). The order of the transfer conditions (i.e., names and practice) was balanced across subjects and the particular devices; procedure and versions were randomly selected for each subject and condition.

Results

The average number of errors for the training and transfer tasks in each condition are shown in Table 2. As in Experiment 1, an unbiased measure was created by summing

Table 2

Raw errors and standard errors for the training and transfer tasks as a function of interface layout, control names and practice

Same Interface Layout

	Training Task		Transfer Task	
	Identical Names	Synonymous Names	Identical Names	Synonymous Names
High-practice 10	8.75 (0.63)	8.2 (0.52)	0.3 (0.11)	2.4 (0.70)
Low-practice 2	8.0 (0.61)	7.8 (0.68)	3.5 (0.69)	5.4 (0.49)

Different Interface Layout

	Training Task		Transfer Task	
	Identical Names	Synonymous Names	Identical Names	Synonymous Names
High-practice 10	9.5 (0.72)	10.2 (0.58)	4.5 (0.77)	7.5 (0.66)
Low-practice 2	9.5 (0.62)	8.6 (0.51)	7.3 (0.90)	9.4 (0.47)

the number of errors each participant made on the first two complete trials of both the training and transfer tasks.

Figure 3 shows the amount of transfer (i.e., log relative errors) as a function of practice, interface and control names. Higher scores indicate more transfer. The pattern of results is similar to the pattern obtained in Experiment 1. More specifically, there appears to be an effect for practice, controls names and interface layout. Furthermore, there also appears to be an interaction between practice and interface, suggesting that the effect of practice is due to indirect transfer. In contrast to the results of Experiment 1, practice and control names in this study were almost perfectly additive. In other words, the pattern also suggests that practice has an effect on indirect transfer. The evidence for these impressions is examined by comparing its fit to several plausible alternatives models. The models used to test the data are the same models used in Experiment 1; therefore the reader can refer to Experiment 1 for the justifications for each model.

Effect of Practice. One possibility is that there is an effect for interface layout and control names but no effect for practice. However, when a model that contains and effects for interface and control names was compared to a model that predicts an effect for all three factors, λ was over 1000, indicating that the data provide stronger support for a model that includes an effect for practice above and over the effects of interface and control names.

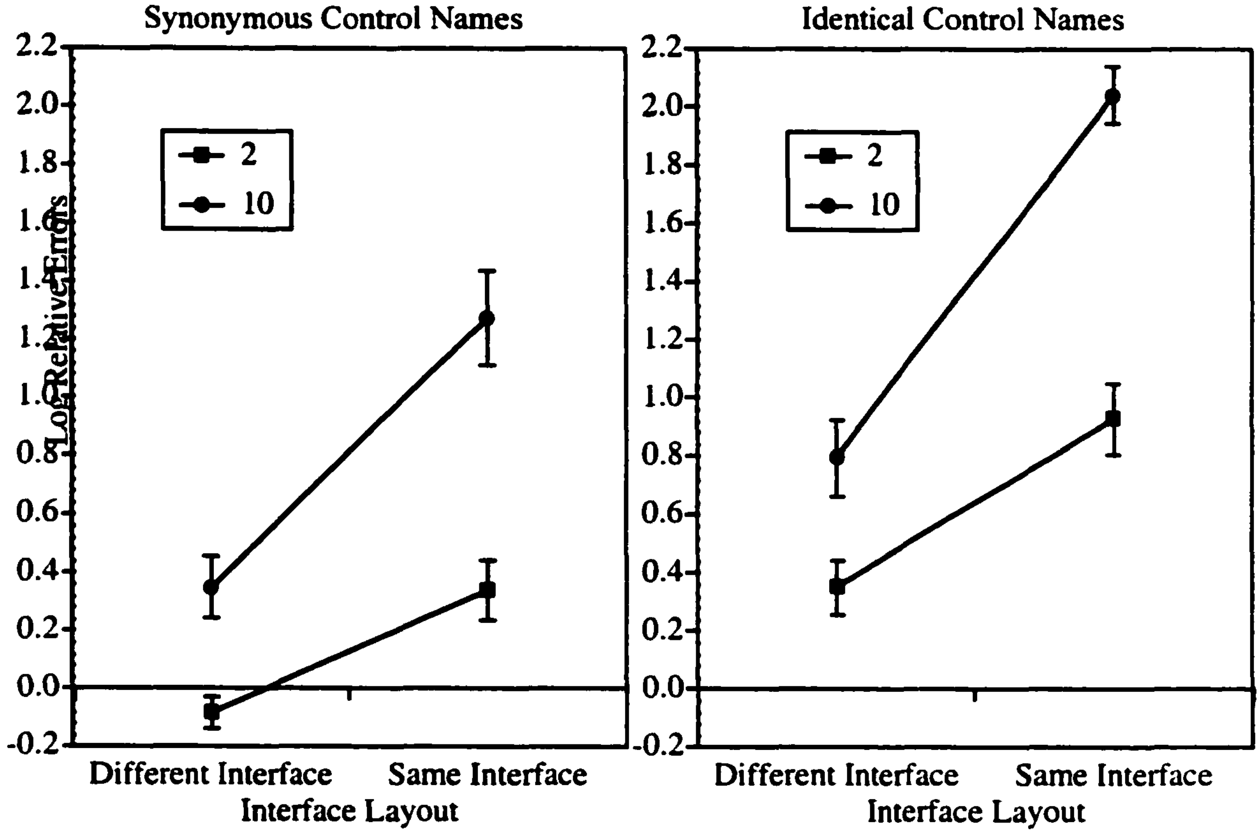

Figure 3. Transfer as a function of practice, interface layout and control names.

Effect for Control Names. Another possibility is that there is an effect for practice and interface layout, but no effect for control names. However, when a model that contains and effects for interface and practice was compared to a model that predicts an effect for all three factors, λ was over 1000, indicating that the data provide stronger support for a model that includes an effect for control names in addition to the effects of interface and control names.

Effect for Interface Layout. A third possibility is that there is an effect for practice and control names, but no effect interface layout. When a model that contains an effect for practice and controls names only is compared to a model that contains an effect for all three factors, λ was over 1000, indicating that the data provide strong support for a model that includes an effect for interface layout names above and over the effects of interface and control names.

Interaction between Interface Layout and Practice. To determine whether practice had an influence on direct or indirect transfer, a model that predicts additive effects for practice controls names and interface layout was compared to a model that also predicts these effects, but in addition, includes an interaction between interface layout and practice. When the simpler additive model was compared to the model that also includes an interaction, λ was 282.99, indicating that the data provide strong support for a model that includes an interaction between interface layout and practice over a model that does not include such an interaction. Thus, it is likely that practice affects the process of indirect transfer and the learning rate.

Interaction between Control Names and Practice. In a similar vein, a comparison was made between a model that included effects for control names, interface layout and practice to a model that included these effects plus an interaction between control names and practice. This comparison was made to determine whether practice had an effect on indirect transfer: If practice was additive with control names then one might expect practice to influence indirect transfer. When the simpler additive model was compared to the model that also includes an interaction, λ was 1.21 indicating that the data provide little support for a model that includes an interaction between control names practice than a model that does not include such an interaction. However, it is useful to evaluate a model that compares an interaction of a particular size, rather than simply test a model that predicts an interaction of any size. As before, an interaction with a size 20% was considered as a theoretically important size of an interaction and this approximately translates into a difference of 0.2 log units. When the additive model was compared was compared to a model that predicts an interaction with a magnitude of .20, the value of λ was over 1000, indicating that the data provide strong evidence for the additive model over the interactive model. Thus, there is strong support for a model based on additive effects of control names and practice compared to a model that also includes a strong interaction between the two. In other words, the data indicate that practice affect indirect transfer and the learning rate.

Discussion

Overall, the results of Experiment 2 replicated the findings of Experiment 1. The superficial manipulations of interface layout and controls names influenced how well the learners could transfer their prior experience to a related device. Furthermore, the replication of the interaction of practice and interface supports the idea that practice affects indirect transfer and the learning rate. The effect of practice was specific to the interface layouts: The effect of practice was much greater when the interface was the same from training to transfer. In contrast to the findings of Experiment 1, there was clear evidence that practice and controls names were additive in Experiment 2. According to the selective effects method, it is likely that practice affects indirect transfer and the learning rate.

However, the additive pattern between practice and control names suggests that the effect of practice was not specific to the control names. Practice improved transfer in both the superficially different condition (synonymous control names) and the identical control names condition. Thus, it is unlikely that an indirect transfer mechanism had an influence on the control names. If practice had an effect on learning the control labels, then one might expect an interaction in which the effects of practice were greater in the identical names condition. In this case, practice would increase the chances of remembering the control label. When the user sees the same label in the transfer phase, they notice that their prior experience was related, and consequently the user could directly transfer their prior experience of the label and step order information to execute the transfer task. When the label is synonymous, the user does not notice that their experience is related, and thus direct transfer does not occur.

Second, if practice increased the chances of noticing the meaning of the controls, then one might also expect an interaction between practice and control names: control names would have an effect in the low practice conditions but not in the high practice conditions (see Figure 4 below). A heightened awareness of the meaning of the control should also increase the chances of noticing the transfer devices are related because the meaning of the controls is also similar. Once this noticing takes place, the user can directly transfer their experience of the steps and their orders. In any event, both interpretations involve interactions of practice and control names, and consequently and direct transfer; however these patterns did not occur. A more parsimonious explanation is that practice had an effect on indirect transfer by influencing a third variable. One possibility is that an increase in practice allows the user to become more familiar with the task goal. A heightened awareness of the task goal would provide a structure to integrate the steps and orders of the transfer task. That is, the goal structure would make the steps and order easier to learn and as a result, an increased learning rate.

Furthermore, this pattern of results makes it difficult to discern whether the effects of practice are specialized or generalized. On the surface, one might suggest that these results mean that practice allows for generalizations because there is an equal

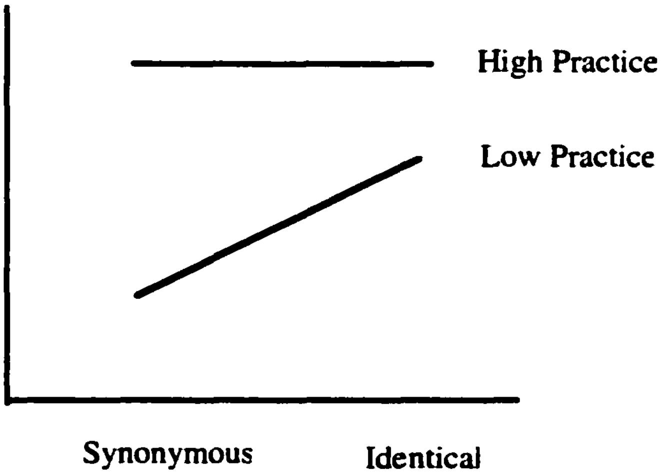

Figure 4. Predictions generated by a pure generalized view of transfer.

benefit for practice in both the superficially similar (identical names) and superficially dissimilar conditions (synonymous names). However, a pure generalized view would predict that there should be no difference between the synonymous and identical control names conditions when practice is high (see Figure 4 below). In other words, one might have expected to find an underadditive interaction between control names and practice. In that case, it is expected that the increased level of practice would allow the user to grasp the meaning of the controls. When they see the synonymous controls of the transfer device, they should be able to notice that their prior experience is related. Once the user notices that their experience is relevant, they could identify the corresponding synonymous controls, and directly transfer their prior experience of the steps and their orders. Nevertheless, this pattern of results did not occur: The effect of control names was the same irrespective of the level of practice.

One explanation for the observed pattern is that practice does in fact promote generalization of the control names (i.e., direct transfer), but the generalizations were blocked because the similarity between the meaning of the synonymous controls was poor. In other words, practice did not increase the transfer of the synonymous condition to the level of the identical condition because the level of similarity in meaning between the synonymous controls was low. For ease of exposition, this hypothesis is referred to as the "poor synonym" hypothesis. If the synonyms used in these experiments shared a low level of meaning, then one might expect that the identical condition would always be superior to the synonymous condition irrespective of the level of practice. This prediction is based on the assumption that practice produces generalization because an increase in practice corresponds to increase in the chance of acquiring the meaning of the controls, and consequently, noticing that prior experience is relevant. Thus, if there is little shared meaning, then there is no simple way that practice can promote transfer in the synonymous condition. In sum, it might be the case that the effects of practice on learning the control names are fully general, but the data fail to show this pattern because there is not enough similarity in meaning between the synonymous controls to allow for generalization. The goal of the third experiment is to test this idea by examining whether the synonyms used in these experiments were actually similar in meaning.

Chapter 4

Experiment 3

The purpose of this experiment is to determine whether poor-synonym explanation can account for why high practice did not promote the complete generalization of the control names. If the control names are not very similar in meaning, then one cannot expect high practice to produce generalizations over the synonymous control names. However, if a high level of similarity in meaning does exist, then there would be the potential for transfer when the names were synonymous. In turn, this would suggest that transfer was lower in the synonymous-names conditions because the user did not fully process the meaning of the controls.

To test for this possibility, 10 participants who had not participated in any of the previous experiments were asked to rate the similarity of words on a five-point scale. In total there were 160 pairs of words. Forty of the word pairs consisted of the synonyms used on the devices (e.g., Gauges - Indicators). Forty other word pairs were created by haphazardly selecting a word from *Random House Roget's Thesaurus*, and pairing it with one of its synonyms (e.g., Banish - Expel). Another 40 pairs were created by randomly re-pairing the 80 device words with each other such that each word was not paired with its synonym (e.g., Gauges - Memory). The final 40 words were created by randomly re-pairing the thesaurus synonyms, taking care to ensure that the pair were not synonyms (e.g., Banish - Banner). Consequently, there were four types of word pairs: Synonymous control names; different control names; synonymous words; and different words. If the control names used in this research are "good" synonyms, then the difference between the synonymous control names and the different control names should be comparable to that obtained for thesaurus synonyms. Conversely, if the control names were poor synonyms, then one might expect the difference between synonymous control names and different control names to be substantially smaller than that obtained for the thesaurus words.

Method

Participants

The participants consisted of 10 students from the University of Alberta who were paid for their participation.

Materials

The participants were provided with a booklet that contained 160 pairs of words: Forty of the word pairs consisted of the synonyms used on the devices in the previous experiments (e.g. Gauges, Indicators). Another 40 pairs were created by randomly pairing the 80 device control names used in experiments with each other such that each control name was not paired with its intended synonym (e.g., Gauges, Memory). Forty other word pairs were created by pseudorandomly selecting a word from *Random House Roget's Thesaurus*, and by pairing it with one of the synonyms provided in the thesaurus (e.g., Banish, Expel). In most cases the experimenter chose the first or second word that was listed as a synonym for the particular word in question. The words chosen from the thesaurus were matched with the control names in terms of the frequency of nouns, and verbs in each list. Thus, if the device controls consisted of 47 nouns and 33 verbs, then the words chosen from the thesaurus also consisted of 47 nouns and 33 verbs. The final 40 words pairs were created by selecting one the 40 words chosen from the thesaurus and randomly pairing it with one of the 40 chosen synonymous of those words taking care to ensure that the pair were not synonyms (e.g., Banish, Banner). As a result, there were four types of word pairs: Synonymous control names; different controls names; synonymous words and different words. The order of the conditions and the particular word pairs were randomized for each participant. Participants were also provided with instructions that explained that the goal of the task: On a five point scale rate how similar the word pairs are in terms of their meaning, where five indicates very similar, and one indicates very dissimilar in meaning. Optical scoring sheets were used to record the responses.

Design and Procedure

The participants were tested individually or in groups of 2 people. Upon arriving, the participants signed a consent form and they were provided with a booklet that contained the 160 word pairs. The participants were asked to read the instructions that explained the rating task. Participants were instructed to rate the word pairs on a five-point scale, where 1 means the words are very dissimilar in meaning and 5 means the word pairs very similar in meaning. After the participant completed the task, they were debriefed and thanked for their participation.

Results and Discussion

A mean rating score was tabulated by averaging the participant's ratings for each condition. One item in the different word condition was eliminated from each participant's score because of a randomization problem. Two other data points were excluded because the participant either left the items blank or they circled more than one response. Figure 5 shows the rating as a function of condition. The standard errors are included in the figure, but they are too small to see. As can be seen from the figure, both the synonymous words and the synonymous control names were rated as being more similar in meaning than either the different words or the different control names. Moreover, the similar ratings for the synonymous words and the synonymous control names suggest that the control names were at least as "good" as the synonyms provided by the thesaurus. Below, these impressions are evaluated by examining the evidence for a model that contains an effect for only the word meaning verses a model that predicts effects for both word type and word meaning plus an interaction.

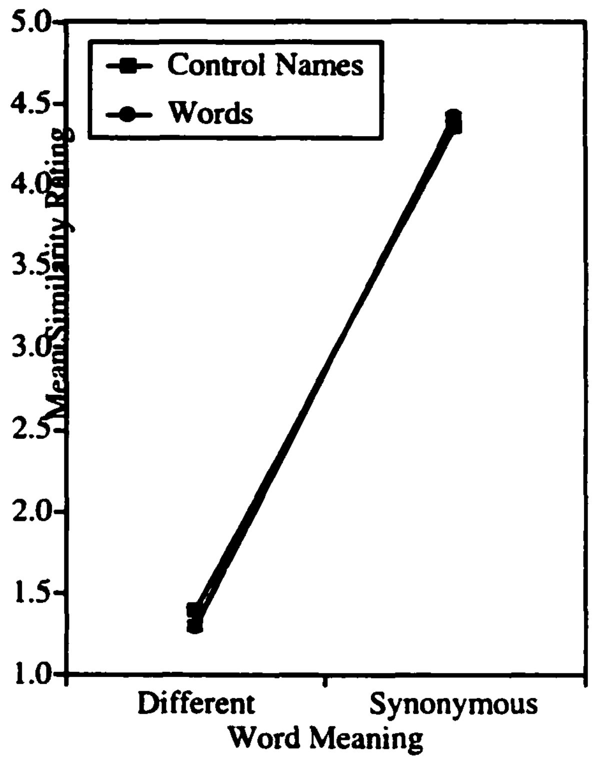

Figure 5. Mean similarity rating as a function word meaning and word type.

Effect for Word Meaning. One possibility is that the control names used in this research are as similar in meaning as the synonyms selected in the thesaurus. Furthermore, if the synonym manipulation was effective, then both the synonymous controls and the synonymous words should be rated as being more similar in meaning than either the different words or the different controls. This model includes an effect for word meaning, no effect for word type, and no interaction between the two. Another possibility is that the control names are rated differently than the word chosen in the thesaurus. More specifically, the poor synonym hypothesis predicts that the synonymous words are rated as being more similar than the synonymous control names. In other words, there should be a smaller effect of word meaning in the synonymous control name condition. This model would include an effect for word type, word meaning, and interaction between the two. However, it is useful to evaluate a model that compares an interaction of a particular size, rather than simply test a model that predicts an interaction of any size. An interaction with a size of 1.0 was considered to be theoretically important. When the additive model was compared was compared to a model that predicts an interaction of this magnitude, the value of λ was over 322.77, indicating that the data provide strong evidence for the simpler model that includes an effect for meaning only compared to a model that includes a substantial interaction.

The results of this experiment indicate that the synonymous controls and the synonymous words were rated as being highly similar in meaning. Moreover, these word pairs were also rated as being more similar than the different words and the different control names. This indicates that the participants were able to perceive the similar meaning of synonyms. Most importantly, the level of rated similarity did not differ between the control names and the words. Thus, the control names used in Experiments 1 and 2 were about as good as synonyms provided by a thesaurus. In other words, it is unlikely that the lower level of transfer for the synonymous high practice condition as compared to the identical high practice condition in the previous experiments was simply as result of a low level of meaning between the synonymous controls names. The results of this experiment suggest that the similar meaning was available to the participants, but they never took full advantage of this possibility. That is, the users never processed the meanings, and therefore they were not able to notice the relations between the training and transfer devices. Without this reminding, the user could not directly transfer their prior experience.

Accordingly, it is more likely that practice affects indirect transfer as was originally proposed. That is, the level of transfer in the high-practice condition was greater than the low practice condition because the practice increases the familiarity of the task goal. This familiarity provides a structure for interpreting and integrating the steps and order information of the transfer device. Such a structure makes the steps and orders of the transfer device easier to learn. When the steps are easier to learn, the rate of leaning the transfer task is higher.

One reason for the finding why practice did not promote the transfer of the synonymous control names to the level of the identical control names is that practice does not promote deep processing. In other words, an increase in practice does not correspond to a large increase in the probability that the user will acquire the meaning of the controls. In this manner, the repetition of a task merely promotes superficial processing, and consequently, the user does not process the meaning of the control to a deep enough level to allow for complete generalization. The next experiment was designed to determine whether the *type* of practice is critical in promoting a deeper level of processing. If this is true, then it should be possible to that one type of practice is more effective than another for promoting transfer to a superficially dissimilar device.

Summary of Experiments 1, 2, and 3.

The first 3 experiments were designed to assess the effect of practice on the transfer across the superficial manipulations of the control names and the interface layout (Experiment 1, and 2). Furthermore, the research was designed to determine whether the additive pattern between the control names and practice in the first two experiments was due to the low level of similarity between the meanings of the synonymous control names (Experiment 3). The investigations revealed that practice improves the acquisition of a novel device through indirect transfer by increasing the rate in which the novel device is learned. My interpretation is that the increased practice made the both the position of the controls and the goal of the procedure more familiar and distinctive to the user. In turn, this level of familiarity and distinctiveness increases the rate in which the device is learned because the goals and the similar interface layouts provide a structure for integrating the step and order information. With a familiar and distinctive structure, the user can more easily encode the steps and order information of the transfer task.

Moreover, the effects of practice were specific to the interface layouts: Practice had a large effect on transferring to an identical interface layout, but a smaller effect on transferring to a different interface layout. This result suggested that the effects of practice on learning the control names are specialized. In contrast, practice facilitated the transfer of the control names to an equal amount: The effect of practice was the same in both the synonymous and identical conditions. These results suggest that practice can promote some generalization of the control names, but total generalization did not occur. I hypothesize that practice did not promote the total generalization of the control names because the type of practice used in the previous experiments failed to promote a deep awareness of the meaning of the controls.

Chapter 5

Experiment 4

The purpose of the fourth and final experiment was to determine whether the type of practice affects how deeply the meaning of controls are processed. One could argue that the form of practice used in the previous experiments (i.e., repetition of the same task on the same device) promotes a relatively superficial level of processing. Increases in this form of practice might not help the user to acquire a complete awareness of the meaning of the controls. Without the awareness of the meaning, the user is unaware that their prior experience is relevant, and consequently, direct transfer cannot occur. In order for transfer in the synonymous condition to increase to the level of transfer in the identical condition, the user must process the meaning of the controls. In this experiment, I evaluate whether practice with multiple examples can facilitate a deeper understanding of the control names, and increase transfer when the names are synonymous.

Prior research indicates that providing multiple examples helps the user to focus on the deep or structural aspects the problem. For example, when a participant compares two superficially different examples that contain the same structure, they develop a schema that contains the common deep and structural aspects of the examples (Gick & Holyoak, 1983). In a similar vein, the presentation of multiple examples of a control name might allow the user to become more aware of the shared meaning of the controls. Consequently, they should be more likely to directly transfer this knowledge when they encounter a device that has a different but synonymous names. With the presentation of two synonymous control names, the user might be able to pay more attention to the meaning, than if they were to simply repeat the same task on the same device over and over again. In this experiment, multiple examples were created by providing the participant with two versions of the same device procedure. Both versions of the devices contained the same task goal, procedure, graphics, gauges and interface layout, but the control names of each version were synonymous rather than identical. Providing the learners with two versions of the control names should help them identify the common semantic content of the names and therefore make it easier to directly transfer this information to a third device with another set of synonymous names. For example, in the training phase the user might learn how to operate one device with the control labeled *Water*, then another with the control labeled *Fluid*. In the transfer phase the learner would be asked to use a device with the control labeled *Liquid*.

The fourth experiment used a 2 x 2 within-subjects design with practice and multiple examples as independent variables. The number of examples was either one or two, and as before, practice was either high or low (2 or 10 trials). As in the previous experiments, there were 4 training tasks, each with its own unique goal and procedure. However, each of the multiple examples of a given task involved the same task goal and procedure. In the low-practice, single-example condition, participants were required to carry out each procedure twice on the same device and procedure. In the low-practice, multiple-examples condition participants were required to complete each of two versions of the procedure once. In the high-practice, single-example condition, the users were required to execute the procedure 10 times on the same device. In the high-practice multiple-condition, participants were asked to operate each version of the procedure 5 times.

For example, in the training phase, the participants in the multiple example condition might operate one device with the control name *Start*, then on the other example they would be required to use a device with the name *On*, while in the transfer phase they would use a device with the control name *Power*. In the single-example case, the learner would operate one device with the control name *Start* in the training phase, and in the transfer phase they would operate a device with the control labeled *Power*. All of the devices in the training and transfer phases (as well as in the multiple example training condition) used the same interface layout. This action was taken to control the impact of the interface layout and to determine whether multiple examples could affect transfer above the influence of the interface layouts.

Method

Participants

The participants consisted of 24 students from the University of Alberta who were paid for their participation.

Materials

The materials were the similar to those used in the previous experiments. Twelve experimental weather devices and one practice device was modeled on a computer screen. For each procedure, there were three versions of the device: All of the versions contained the same gauge, graphics and appearance, but the control names in each version were synonymous (e.g., *Temperature; Thermal; Warmth Control*). However the interface between the training and transfer devices was always the same. The instructions were the same as those used in the previous experiments except that the user was told that they might be required to operate a given device more than once.

Design and Procedure

The procedure was the same as in the Experiments 1 and 2. Participants first operated four training devices, and then they were required to operate four transfer devices. The independent variables were practice and the number of examples. In the low-practice, single-example condition, the participants completed the same procedure on the same version of the device twice (i.e., identical control names). In the low-practice, multiple-example condition, the user executed the procedure twice: once with one version and once on the other version. In the high-practice, single-example condition participants operated the device procedure 10 times on the same device. In the high-practice, multiple-example condition the user also executed the procedure 10 times, but 5 times with one version and five times with the other version. The versions of multiple examples and the procedures were randomly selected for each participant, and order of the conditions was balanced across subjects.

Results and Discussion

The transfer measure used in the first two experiments was calculated by summing the total number of errors the participants made on the first two trials of the transfer task and dividing this by the sum of the errors made on the two trials of the transfer task. Using this measure in the current design was problematic. The first two trials of the low-practice, multiple-example condition (and in some cases for the high-practice, multiple-example condition) involved two different versions of the same procedure, while the two trials in the single-example conditions always involved only a single version. To avoid the potential bias, only the errors made on the first training trial were used in the calculation of the training errors. As in the previous experiments, the transfer errors was calculated by summing the total number of errors the participants made on the first two trials of the transfer task. The overall transfer measure was obtained in the same manner as the previous experiments: Transfer = ln [training errors +1]/[transfer errors +1]). The average number of errors for the training and transfer tasks in each condition are shown in Table 3.

Table 3

Raw errors and standard errors for the training and transfer tasks as a function of practice and examples

	Training Task		Transfer Task	
	Single-example	Multiple Example	Single-example	Multiple Example
High-practice 10	7.0 (0.29)	6.8 (0.32)	3.3 (0.55)	1.5 (0.45)
Low-practice 2	7.2 (0.33)	7.1 (0.37)	4.7 (0.50)	4.7 (0.42)

Figure 6 shows the amount of transfer as a function of practice and interface layout. As can be seen from the figure, multiple examples only has an effect on transfer when the practice is high; there was no benefit for providing multiple examples when practice was low. A priori, one might predict an effect for both practice and multiples:

Previous work indicates that both practice (e.g., Jolicoeur & Milliken, 1989) and the provision of multiple examples (e.g., Ross & Kennedy, 1990) can facilitate the transfer of prior experience. Below, an additive model that predicts an effect for both practice and multiple examples is compared to several alternatives for that data.

Effect for Multiple Examples. One alternative to the additive model includes an effect for practice, but no effect for multiple examples. An effect for practice is predicted because the findings of Experiment 1 and 2 indicate that practice helps user to improve transfer across the control names. In contrast, no effect for multiple examples might be predicted because the interface between the training and transfer devices was always the same. With this form of similarity, participants might quickly learn to ignore the control names and operate the devices by locating the positions of the buttons. When the additive model is compared to the model that predicts a single effect of practice λ was over 1000, indicating that there was strong evidence for a model that includes an effect for multiple examples.

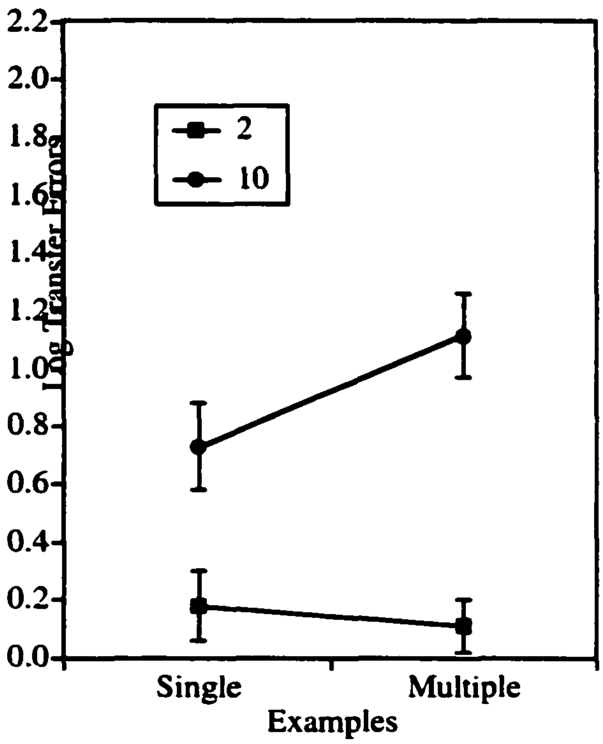

Figure 6. Transfer as a function of examples and practice.

Effect for Multiple Examples only in High-practice Condition. Another possibility is that there is a specific interaction between multiple examples and practice, with an effect for multiple examples only in the high-practice conditions and no effect of examples in the low-practice conditions. According to this logic, the user doesn't notice the similar meaning of the control names with only a single presentation of each synonym; access to this information only occurs after several exposures to the names. When this model was compared to a model that predicts an effect of practice only λ was 20.86 indicating the data support a model that predicts a specific interaction between examples and practice over a model that predicts an effect of practice only.

The results of Experiment 4 indicated that the provision of multiple examples could promote the transfer of prior experience to superficially dissimilar devices. This result also shows that the effect of multiple examples occurs above and beyond the effect of the interface layout: There was an effect of multiple examples even though the interface layout was always the same from training to transfer. Thus, the argument that the user's operated the devices on the bases of the interface alone cannot be true. Multiple examples can influence transfer independent of a similar interface layout.

Presumably, the provision of multiple examples allows the user to become more aware of or pay more attention to the meaning of the controls. With this increased sense of the meaning, the user is able to notice that their prior experience is relevant to the transfer task. Accordingly, this noticing can allow the user to map the meanings of the corresponding controls and directly transfer their knowledge of the step order to execute the transfer task. However, the results reveal that practice and examples interact when the learner is required to operate a device with synonymous control names. That is, multiple examples of a control name allow the learner to generalize their knowledge over the synonymous control names only when practice is high. When practice is sufficiently low, the users are unable to take advantage of the multiple examples. Thus, while the provision of multiple examples increases the participant's awareness for the meaning of the controls, this awareness or attention is only facilitated with extended practice. Presumably, a single exposure to each control name is not sufficient to tap into the control name's meaning. Hence, the conclusion drawn from this experiment is that the type of practice is important in determining whether the learner is likely to directly transfer their prior experience to a superficially dissimilar device. Simple repetition of the same task on the same device promotes only superficial processing of the control names, and as a result, transfer can only be facilitated through indirect means. In contrast, the repetition of a task on multiple examples of a device provides deeper processing of the meaning of the controls, and therefore direct transfer.

Chapter 6

General Discussion

The current investigation was designed to address four issues. First, the research was intended to establish whether a simple learning strategy such as practice could enable users to transfer their prior experience to a structurally similar but superficially dissimilar task. Second, I hoped to evaluate whether the effects of practice were specialized or generalized. The specialized view asserts that practice enables the user to optimally execute the practiced task, but transfer to related tasks is limited. In contrast, the generalized view asserts that practice can allow the user to develop generalizations of the task, and consequently, practice should enable the user to transfer their experience to associated tasks. Third, the present work was designed to ascertain how practice was related to the transfer of device experience. More specifically, the experiments were designed to determine whether practice influenced direct transfer and the initial error rate, or indirect transfer and the learning rate. Finally, I wanted to determine if the type of practice was important in determining whether transfer would occur to superficially dissimilar domains.

With respect to the first issue, the current findings suggest that practice can allow the user to transfer their prior experience to a superficially dissimilar device, but the amount of transfer depends on the type of practice and the form of superficial similarity. The results of Experiment 2 indicate that practice with a single device can indirectly facilitate the transfer of across the synonymous controls names. However, the benefit of this form of practice did not allow for the complete generalization of the control names: The effect of practice did not raise the level of transfer in the synonymous condition to the level of the identical condition (see Figure 4). In Experiment 4 it was found that the type of practice was critical in determining whether the user would directly transfer their experience to a device with superficially dissimilar control names. When users are given multiple examples of a procedure, they more likely to display transfer to a device with synonymous control names than when they are given only a single example. Moreover, this effect only occurred when the level of practice was high. In contrast, the results of

Experiment 1 and 2 indicate that practice was specific to the interface layout: Practice greatly improved the transfer in the identical interface condition, but it had only a minor influence learning a device with a different interface.

The second aim of the study was to determine whether the effects of practice are generalized or specialized. The results of the current experiments suggest that the answer to this question also depends on the type of practice. When users practice a task on the same device (single example), the ability of practice to improve the acquisition of superficially dissimilar devices (different interface, synonymous control names) is limited. Transfer in the both the high-practice, synonymous-control-name and the high-practice, different-interface conditions never reached the level of transfer in the high-practice, identical-name and the high-practice, same-interface conditions. That is, high-practice never improved performance with the superficial dissimilar devices to the level of the identical or similar devices. While practice with a single example improved the transfer in both the synonymous and identical conditions (suggesting some generalization), the facilitation was not fully general, because transfer in the identical names condition was always higher than the synonymous condition irrespective of high practice. Thus, these results suggest that the effects of practice with a single example are mainly specialized.

However, when users were provided with high practice and multiple examples of the control names, they were able to generalize this knowledge to a device with synonymous controls (direct transfer). These results suggest that the effects of varied practice (high-practice with multiple examples) promotes generalization. Thus, the findings imply that the effects of practice are mainly specialized when one practices the same task on a single version of a device, but the effect of practice are general when one practices a task on multiple versions of a device. While the effects of multiple examples and interface were not examined in this study, I suspect the results would be the same: Multiple examples would promote the generalization of the interface layouts. This is a direction for future research.

In terms of the third aim, the results indicate practice with a single example had an effect on indirect transfer and the learning rate. In Experiment 2 practice was additive with the control names manipulation. Prior work has indicated that the control names influence direct transfer and the initial error rate (O'Reilly & Dixon, 2001). According to the logic of the selective effects method, any variable that is additive with a variable known to affect direct transfer (e.g., control names) may have its effect on indirect transfer. In a similar vein, practice also interacted with a variable hypothesized to influence indirect transfer, interface layout, in Experiments 1 and 2. According to the selective effects method, any variable that interacts with a variable known to affect indirect transfer should also influence indirect transfer and the learning rate. The interaction between practice and interface and the additive pattern between practice and control names suggests that practice with a single example affects indirect transfer and the learning rate.

In contrast, the findings of Experiment 4 indicate that the effect of practice on multiple examples is a different form of practice than practice with single examples. When a user practices with multiple examples, they are more likely to notice the meaning of the control names. This increased sense of meaning allows the user to notice that the transfer controls, and consequently the transfer device, is related to the training device. With this reminding, the user directly transfers their knowledge of the training steps and order of the steps to operate the transfer device. When the meaning of the control names is available, transfer is simply a matter of mapping the corresponding controls of both devices on the bases of their shared meaning. Once this mapping is complete, operating the transfer is easy because the learner can use the step order of the training device to operate the transfer machine. With respect to the final objective, the results of Experiment 4 indicated that the provision of multiple examples could help the user to transfer their prior experience to a superficially different device. Thus, the type of practice is critical in determine how effective practice will promote transfer.

While the above findings have shed some light on the role of practice on the transfer of device knowledge, several questions remain. First, why does practice with a single example interact with interface and not with the control names? Second, why is transfer in the identical name conditions always higher than in the synonymous control name conditions? Third, how does practice increase the transfer of prior experience? Finally, how does the provision of multiple examples help the user to transfer to a superficially dissimilar device? Below, each of these issues is discussed.

The results of experiments indicate that practice interacts with the interface layout, but it is additive with the control names. The effect of practice is large in the same-layout condition interface because practice makes the locations and relations between the controls more familiar and distinctive. When the user encounters a device with the same configuration of controls, they learn the device faster because they use the familiar configuration as a structure to integrate the step and order information of the transfer task. That is, the same interface layout makes the transfer device easier to learn. Practice has a large effect in this case because there is overlap in the positioning of the controls. However, when the layout of controls is different, the prior experience cannot help because the positions are no longer familiar, and therefore there is no familiar layout structure to help the user to learn the steps and step-orders of the transfer task. In contrast, practice with a single example improves indirect transfer of the control names because practice increases the familiarity of the goal of the procedure. This familiarity provides a structure that makes it easier to learn the steps and the orders of the transfer device, thus the rate of learning is increased. The effect of practice on the control names is additive because the familiarity of the procedural goal occurs in both the identical and the synonymous name conditions.

With respect to the second question, previous work has provided several reasons why transfer is typically greater in the identical control name condition than the synonymous condition (O'Reilly & Dixon, 2001). When the names are identical, the user notices that the devices are related, and consequently they directly transfer their prior knowledge of the steps and the step order to execute the transfer device. That is, the identical labels remind the user that their prior experience is relevant. In contrast, when the control names are synonymous, the user does not notice that the devices are similar, and as a result they do not use this prior experience to carry out the transfer task. That is, when the names are synonymous, direct transfer does not occur (unless one is provided with multiple examples). As a result, transfer is lower in the synonymous condition than in the identical condition. Moreover, transfer might be greater in the identical control condition because the names are easier to remember than synonymous. In fact, O'Reilly and Dixon (2001b) found that the memorably of the controls can have a large effect on the transfer of device experience. Finally, another possibility is that transfer is poor in the synonymous conditions because the words used in the study were "poor" synonyms. In other words, the users were unable to identify the common meaning of the controls because the synonymous were only remotely similar in meaning. The results of Experiment 3 cast doubt on this hypothesis: The control names were rated as having similar meanings and the ratings did not differ substantially from the similarity ratings of synonyms chosen from a thesaurus.

Another issue is how practice can help the user to increase the amount of transfer to a novel device. There are three categories of explanations that describe how practice can affect performance: Procedure transformation, strengthening, and strategy selection (Lee & Anderson, 2001). Proponents of the procedure transformation view assert that a task is composed of multi-step procedures. Practice is thought to improve performance by either chunking (Newwell and Rosenbloom, 1981) or compiling (Neves & Anderson, 1981) the multi-step task into a macro procedure that contains fewer steps than the original procedure. However, these explanations cannot account for the effect of practice in simpler tasks that do not have multiple steps. According to strengthening theories, practice improves performance by increasing the probability that a particular node or other form of knowledge representation is activated (e.g., Anderson, 1982). With increasing practice, the nodes that represent information are more likely to become active, and, as a result, the appropriate response is faster and more probable. Finally, selection theories posit that users have several methods for executing the task and that each method varies in terms of how long it takes to complete the procedure. Participants become quicker at executing the task because they select faster methods to carrying out the procedure (e.g., Crossman, 1959; Logan, 1988).

On the one hand, a proponent of the procedure transformation view can account for the current findings if one assumes that the user complies or chucks the series of steps into a macro procedure. In this manner, the learner performs the task much faster and more accurately by reducing the procedure to a single step. Thus, when a user attempts to carry out the Start Up Procedure, he or she automatically carries out the steps to reach the goal. Transfer in the synonymous control name and the different interface conditions was lower than in the identical conditions because it is assumed that the macro procedure is only executed when the names are identical and when the interface is the same. In other words the instantiation of the macro procedure is dependent upon the superficial relation of the device. However, this view would have difficulty in explaining the effect of practice when the control names are synonymous.

Alternatively, the procedure transformation view cannot account for the results if one assumes that the user collapses the name and interface layout information. A proponent of this view would assert that the representation of the task consists of the control names, the step order and the position of the controls. When practice is low, the learner uses the both the control names and the interface layouts to operate the device. That is, they remember the procedure is terms of the control name and its relative position on the device. However, with increases in practice the user collapses the representation across the name information, and they simply operate the device based on the position of the controls. When the interface is the same, transfer should be the higher than when it is different because the user is simply relying on the location of the controls to execute the procedure. However, this view would have difficulty in explaining why the control names still had an effect even when practice is high. More importantly, the model would not be able to predict why multiple examples of control names facilitate transfer even when the interface is the same (Experiment 4). In this case the model predicts that names (multiple examples) should have no effect when practice is high because the user executes the task on the bases of the positions of the controls only. Thus it is unlikely that the procedure transformation view can account for the results of these experiments if one assumes that the user collapses across both the step and control names.

A strengthening model could account for the both the interface and the name results. Practice has an effect in the similar interface conditions because the nodes that represent the spatial locations of the controls are facilitated through practice. Increases in practice would increase the chances of activating the memory for the spatial locations of the controls. The effect of practice in the different interface conditions is lower because there was no overlap between the stored training representation and the novel representation. In a similar vein, transfer is higher in the identical control name condition because the representation of the transfer controls is facilitated by the strengthened training representation. However, the model has difficulty in explaining the facilitation effect of practice on the synonymous controls names. In this case, practice has no way of increasing activation because the synonymous words were not directly activated. A strength model can only account for the results if one either assumes that the activation of a word also activates a separate "meaning node", or if the activation of the word also activates the semantic neighbors through spreading activation. If this were the case, transfer should decrease as the semantic distance between the training and transfer words increases. Consequently, the model would predict an effect for practice in the synonymous name condition because practice activates the semantic neighbors of a particular word, or in this case the particular control names. Transfer is greater in the identical case because practice increases the activation of the identical name more than the synonymous name. This is because it is assumed that the strength of the activation decreases as the distance between the activated node and the semantically related nodes increases. The model would predict an effect for multiple examples because with high practice, the activation of two synonymous names produces even more spreading activation to the semantically related neighbors. The increased activation of the semantic neighbors should increase the probability of activating any other name that is semantically related to the multiple examples.

At first glance, the strategic selection explanation might have some difficultly in explaining the results. In the current experiments, there are two possible ways to execute the task: Remembering the sequence in terms of the positions of the controls, and remembering the sequence in terms of the names of the controls. Evidently, prior work indicated that both the position (interface) and the control names influence transfer at the same magnitude (O'Reilly & Dixon, 2001). Thus, switching from one method to the other should not increase transfer, and consequently method switching should not affect transfer. However, it could be the case that participants are switching from a name-based strategy to a position based strategy with increases in practice (with a single example). Recall that the practice manipulation in Experiments 1 and 2 required the learners to execute the task on the identical training device. With increased presentation, the user might ignore the constant presentation of the identical names and simply carryout the task with respect to the positions of the controls. If this were the case, then one would expect to observe more of an effect of practice when the interface was the same; however, when the interface is different, transfer should be greatly reduced. In fact, this is the pattern of results observed. On the other hand, when the training conditions are varied (multiple examples) the user does not ignore the control names because the training devices are slightly different and changing. Consequently, with increases in practice, the participant pays attention to both the positions and the names. This hypothesis predicts that an effect for practice on the control names. In fact the data in Experiment 4 indicate that practice does have an effect on the transfer across the control names but only when the practice is high. However, this strategy cannot account for the effect of control names during practice with a single example. In that case it is expected that the user would execute the procedure on the bases of the interface layouts and not the control names. However, the results indicted that control names had an effect when practice was carried out on a single example. Thus, the strategy explanation cannot account for the current findings.

Although the strengthening explanation can account for the results, I prefer to explain the results according to the device transfer model of O'Reilly and Dixon (2001). According to that model, there are two forms of transfer and two components of learning. When the user has knowledge that is immediately applicable to the task, the learner can reduce their initial error rate substantially by directly transferring this experience to the novel machine. Alternatively, indirect transfer can occur even when the user has no prior bases for selecting one action over another. The steps and the step order of a task can be learned more easily when the user has a familiar and distinctive structure to incorporate the novel information. This in turn increases the rate in which the transfer task is acquired. The interactive and additive patterns of the first two experiments suggest that practice with a single example supports transfer through indirect means. Practice helps the user to transfer their knowledge by providing a familiar structure to integrate the transfer task. This allows the user to learn the transfer task more easily. The model predicts that practice should facilitate the transfer across identical controls names because the identical label allows the user to notice that the training and transfer tasked are related; consequently they can directly transfer their prior experience. However, transfer is lower in the synonymous condition because the user dose not notice that their prior experience is relevant, and therefore, they do not directly transfer their prior experience. When the interface is the same indirect transfer is facilitated because the user has a familiar structure to help them learn the transfer task. Conversely, practice should have little effect on the transfer of a different interface because there is no structure to help learn the transfer task.

The final issue concerns the question of how multiple examples increase transfer. According to the results of Experiment 4, practice with multiple examples influences direct transfer and a reduction of the initial error rate. High practice with multiple examples allows the user to process or pay attention to the meaning of the controls. With this meaning in mind, the user is likely to notice that the training and the transfer devices are related. Once this relation is noticed, the user directly transfers their prior experience of the steps and the step order to the transfer device. The process entails that the learner maps the controls of the training device to the transfer controls on the bases of the similar meaning of the controls. Once this mapping is made, the learner executes the transfer task by using the step order learned in the training task. However, this effect was only evident when the user was provided with high-practice. Thus, the single presentation of two examples is not enough to foster spontaneous transfer to a novel device. The user requires multiple opportunities to carry out the procedures before they are able to process the critical structure of the controls: the semantics. On balance, these results are encouraging because to the best of the author's knowledge, this is the only study to demonstrate that learner's are capable of spontaneously transferring their knowledge to a superficially dissimilar domain with only the aid of multiple examples. In other work, the participants required a hint to spontaneously use this information (e.g., Catrambone &

Holyoak, 1989; Ross & Kennedy, 1990) or an instruction to compare the examples (Catrambone & Holyoak, 1989).

In sum, the role of prior experience on learning a new device can be outlined in the following manner. After the user has been given low practice with executing a task on a single device, their representation of the task is weak. That is, they do not have a solid (error-free) representation of the procedure in terms of both their memory for the controls, and the step orders. Furthermore, the learners do not normally process the meaning of the control names to a deep level. In fact, their memory of control names is mainly centered on the verbatim control labels. The learner only occasionally encodes the verbatim control labels and the step order information, while leaving the semantics of the name relatively uninterpreted. When the user attempts to operate a device with identical control names, they notice that the training and transfer machines are related, and consequently, they directly transfer whatever steps and order information that they can remember. In contrast, when the names are synonymous, the user does not notice that their prior experience is relevant to transfer task, and thus they do not use their prior experience to execute the task. In other words, when the names are synonymous direct transfer does not occur. The implicit assumption of this framework is the control name determines whether the user will notice that the devices are related. When the names are identical, the learners notice the relation because their memory of the training device contains the same verbatim label as the transfer controls. In contrast when the name is synonymous, the user does not notice the relation because their representation of the training task does not contain the synonymous control label. Because the user does not process the meaning of the controls very deeply, they have no way of relating the devices, and consequently, they fail to directly transfer their prior experience.

On the other hand, high practice on a single example of a device improves performance by increasing the rate in which the transfer task is learned (indirect transfer). Thus, transfer is greater in the same-interface condition because practice makes the configuration and relations of the controls even more familiar than with low-practice. With this familiar structure at hand, the user can learn the transfer device easier by integrating the step and order information of the training task with the familiar interface layout. This improves the rate of acquisition by making the steps of the transfer task easier to learn. The effect of practice in the different-interface condition is low because the interface configuration is no longer familiar, so the user does not have a structure to help them learn the transfer task. The small effect of practice in this condition could be the result of practice increasing the participant's ability to learn the goal of the procedure. This would provide a structure to allow the user to learn the steps of the transfer task more easily.

In a similar vein, practice improves the performance in the name conditions because practice makes the goal of the procedure more familiar. This provides a structure and organization for learning the steps of the transfer device. Thus, the transfer task is learned faster because the goal structure enables the transfer task to be learned more easily. The effect of practice and names is additive because the familiarity with the goal of the procedure occurs in both the identical and synonymous conditions. Consequently, practicing a task on a single example of a device promotes relatively superficial processing of the task. In line with this interpretation is the result that practice has its largest effect on the transfer to a device with the same interface conditions, while it has only a small effect on transferring to a device with a different interface. Accordingly, practice with a single example does not promote the complete generalization of the control names: Practice does not increase the level of transfer in the synonymous condition to the level of transfer in the identical condition. If practice with a single example promoted deep processing, then this difference would not be expected. In other words practice with a single device does not increase the user's awareness of the meaning of the controls.

Finally, practice with multiple examples improves direct transfer by reducing the learner's initial error rate on the transfer task. When practice is high, the user is more likely to pay attention or process the meaning of the controls. Switching from using one device with a particular set control names to an identical device with a different (but synonymous) set of names should cause the user to notice or pay attention to the names. For instance, Ross and Kilbane (1997) found that participants were better able to transfer their prior experiences of a superficially dissimilar task when they were trained with similar examples that differed in only one respect. Ross and Kilbane hypothesized that the provision of similar examples allowed the participant to notice the single dissimilar aspect of the example because the similarity of the example highlighted the superficial difference. As a result, the learner was more likely to develop generalizations, and they were less likely to be affected by superficial dissimilarities. The experimenters found evidence for this notion. Thus, providing the users with devices that are similar but differ only in the control names should allow them to notice the different control names. When they take notice of the names, they are more likely to notice the similar meanings of the controls. In turn, this enables the user to notice that the training and transfer devices are related. Once this noticing has taken place, the learner can directly transfer their prior experience to the transfer task by mapping the corresponding controls and step order information.

However, when there is one switch (low-practice) the user might notice that the names were changed, but they might simply assume that the names, and consequently the devices, were totally different. In this case the user will not directly transfer their prior experience. Conversely, when users are continually given alternating versions of the names (high-practice), they may be more likely to notice that the names are similar in meaning. This would increase the chances of noticing that the devices were similar, and as a result, the user would directly transfer their prior experience.

The moral of this research is that that practice with multiple examples promotes deep processing, and that this experience can be effectively used to operate superficially dissimilar devices. However, practice with a single example primarily facilitates superficial processing, and, as a result, transfer to a superficially dissimilar device is lower than if one were provided with varied practice.

Conclusion

The proponents of the doctrine of formal discipline asserted that transfer amongst various domains was possible through the study of classics. In contrast, Thorndike argued that transfer across different domains was constrained by the number of identical elements contained in the tasks. However, Wertheimer advocated that the most effective way to foster transfer was provide the user with varied training conditions. The current findings suggest that transfer is less flexible than the proponents of formal discipline believed, but transfer is also more flexible than Thorndike claimed. On balance, one of the most effective ways to foster transfer is provide the learner with varied training experiences as Wertheimer advocated in 1959. That is, the type of practice is critical in determining how effective a learner will be able to apply their prior experience to a superficially dissimilar domain.

Bibliography

Ahn, W., Brewer, W., Mooney, R. (1992). Schema acquisition from a single example. Journal of Experimental Psychology: Learning, Memory, & Cognition, 18, 391-412.

Anastasi, A., McKeachie, J., Smith, M. & Cummings, N. (1992). A century of psychological science. American Psychologist. 47, 842-843.

Anderson, J. (1987). Skill acquisition: Compilation of weak-method problem solutions. Psychological Review, 94, 192-210.

Anderson, J. R. (1995). Learning and memory. New York:Wiley.

Anderson, J. R., & Fincham, J. M. (1994). Acquisition of procedural Skills from examples. Journal of Experimental Psychology: Learning, Memory and Cognition,20, 1322-1340.

Anderson, J. R, Fincham, J. M. & Douglass, S. (1997). The role of examples and rules in the acquisition of a cognitive skill. Journal of Experimental Psychology: Learning, Memory, & Cognition, 23, 932-945

Angell J. (1908). The doctrine of formal discipline in the light of the principles of general psychology. Educational Review, 36, 1-14.

Atkinson, R.C., Bower, G. H. & Crothers, E. J. (1965). *An introduction to mathematical learning theory.* New York: Wiley.

Bibby, P. & Payne, S. (1993). Internalization and the use specificity of device knowledge. Human-Computer Interaction, 8, 23-56.

Bibby, P. & Payne, S. (1996). Instruction and practice in learning to use a device. Cognitive Science, 20, 539-578.

Blessing, S. B., & Ross B, H. (1996). Content effects in problem categorization and problem solving. Journal of Experimental Psychology: Learning , Memory and Cognition,22, 792-810.

Borgman, C. (1999). The user's mental model of an information retrieval System: an experiment on a prototype online catalog. International Journal of Human-Computer Studies, 51, 435-452.

Bower, G.H. (1961). Application of a model to paired-associate learning. Psychometrika, 26, 255-280.

Butterfield, E., Slocum, T. & Nelson, G. (1993). Cognitive and behavioral analyses of teaching and transfer: Are they different. In D. K. Detterman & R. J. Sternberg (Eds.), Transfer on Trial: Intelligence, cognition, and Instruction (pp.192-257). Norwood, NJ: Albex.

Carraher, T., Carraher, D. & Schliemann, A. (1985). Mathematics in the streets and schools. British Journal of Developmental Psychology, 3, 21-29.

Catrambone, R. (1995). Aiding subgoal Learning: Effects on transfer. Journal of Educational Psychology,87, 5-17.

Catrambone, R. (1996). Generalizing solution procedures learned from examples. Journal of Experimental Psychology: Learning, Memory and Cognition, 22, 1020- 1031.

Catrambone, R., Holyoak, K. J. (1989). Overcoming contextual limitations on problem-solving transfer. Journal of Experimental Psychology: Learning , Memory and Cognition,15, 1147-1156.

Chase, W. & Ericsson, A. (1982). Skill and working memory. In G.H. Bower (Ed.), The Psychology of Learning and Motivation, (Vol. 16, pp. 1-58). San Diego, CA: Academic Press

Cheng, P., Holyoak, K., Nisbett, R. & Oliver, L. (1986). Pragmatic verses syntactic approaches to training deductive reasoning. Cognitive Psychology, 18, 293-328.

Chi, M. T. H., Bassok, M., Lewis, M. W., Reimann, P., & Glaser, R. (1989). Self-explanations: How students study and use examples in learning to solve problems. Cognitive Science, 13, 145-182.

Chi, M. T. H., De Leeuw, N., Chiu, M., & LaVancher, C. (1994). Eliciting self-explanations improves understanding. Cognitive Science, 18, 439-477.

Chi, M., Feltovich, P., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. Cognitive Science, 5, 121-152.

Cohen, J. (1994). The Earth is round ($p<.05$). American Psychologist, 49, 997-1003.

Cox, B. (1997). The rediscovery of the active learner in adaptive contexts: A developmental-historical analysis of transfer of training. Educational Psychologist, 32, 41-55.

Crossman, E. (1959). A theory of the acquisition speed-skill. Ergonomics, 2 , 153-166.

Dixon, P. (1998). Why scientists value *p* values. Psychonomic Bulletin & Review, 5, 390-396.

Dixon, P. & Gabrys, G. (1991). Learning to operate complex devices: Effects of conceptual and operational similarity. Human Factors, 33, 103-120.

Dixon, P., & O'Reilly, T. (1999). Scientific verses statistical inference. Canadian Journal Of Experimental Psychology,53, 133-149.

Dixon, P., Zimmerman, C., & Neary, S. (1997). Prior experience and complex procedures. Memory & Cognition, 25, 381-394.

Ericsson, A. (2001). The path to expert golf performance: Insights from the masters on how to improve performance by deliberate practice. In P.R. Thomas (Ed.), Optimizing Performance in Golf (pp. 1-57). Brisbane, Australia: Australian Academic Press.

Ericsson, A., Krampe, R. & Tesch-Romer, C. (1993). The role of deliberate practice in the acquisition of expert performance. Psychological Review, 100, 363-406.

Ericsson, A. & Polson, P. (1988). An experimental analysis of the mechanisms of a memory skill. Journal of Experimental Psychology: Learning Memory and Cognition, 14, 305-316.

Fitts, M. (1964). Perceptual-motor skill learning. In Categories of Human Learning, ed. A.W. Melton, pp 243-285. New York: Academic

Forbus, K., Gentner, D. (1994). Mac/Fac: A model of similarity-based retrieval. Cognitive Science,19, 141-205.

Fong, G., Krantz, D. & Nisbett, R. (1986). The effects of statistical thinking training on thinking about everyday problems. Cognitive Psychology, 18, 253-292.

Gardner, R. (1990). When children and adults do not use learning strategies: Toward a theory of settings. Review of Educational Research, 60, 517-529.

Gentner, D. & Markman, A. (1997). Structure mapping in analogy and similarity. American Psychologist, 52, 45-56.

Gick, M. L., & Holyoak, K., J. (1980). Analogical problem solving. Cognitive Psychology,12, 306-355.

Gick, M. & Holyoak, K. (1983). Schema induction and analogical transfer. Cognitive Psychology, 15, 1- 38.

Gitomer, D. (1988). Individual differences in technical troubleshooting. Human Performance, 1, 111-131.

Hanish, K., Kramer, A. & Hulin, C. (1991). Cognitive representations, control, and understanding of complex systems A field study focusing on components of user's mental Models and expert/novice differences. Ergonomics, 34, 1129-1145.

Harvey, L., & Anderson, J. (1996). Transfer of declarative knowledge in complex information-processing domains. Human-Computer Interaction,11, 69-96.

Heathcote, A. Brown, S. & Mewhort, D. (2000). The power law repealed: The case for an exponential law of practice. Psychonomic Bulletin & Review, 7, 185-207.

Heydenbluth, C., Hesse, F. (1996). Impact of superficial similarity in the application phase of analogical problem solving. American Journal of Psychology, 109, 37-57.

Hintzman, D, (1986). " Schema abstraction" in a multiple-trace memory model. Psychological Review, 93, 411-428.

Holyoak, K. & Koh, K.(1987). Surface and structural similarity in analogical transfer. Memory and Cognition,15, 332-340.

Howie, D. & Vicente, K. (1998). Making the most of ecological interface design: The role of self-explanation. International Journal of Human Computer Studies, 49, 651-674.

Jolicoeur, P. & Landau, M. (1984). Effects of orientation on the identification of disorientated objects. Canadian Journal of Psychology, 38, 80-93.

Jolicoeur, P. & Milliken, B. (1989). Identification of disoriented objects: Effects of context of prior presentation. Journal of Experimental Psychology: Learning, Memory and Cognition, 15, 200-210.

Jolicoeur, P. (1985). The time to name disorientated natural objects. Memory & Cognition, 13, 289-303.

Kamouri, A. L., Kamouri, J., & Smith, K, (1986). Training by exploration: facilitating the transfer of procedural knowledge through analogical reasoning. International Journal of Man and Machine Studies, 24, 171-192.

Kieras, D. E., & Bovair, S. (1984). The role of a mental model in learning to operate a device. Cognitive Science, 8, 255-273.

Kitajima, M. & Polson, P. (1997). A comprehension-based model of exploration. Human-Computer Interaction, 12, 345-389.

Koehler, W. (1927). The mentality of apes. (2nd Rev. ed., E. Winter, Trans.) New York: Harcourt Brace

Kolers, (1976). Reading a year later. Journal of Experimental Psychology: Human Learning and Memory, 2, 554-565.

Kolers, P., & Perkins,D. (1969). Orientation of letters and their speed of recognition. Perception and Psychophysics, 37, 429-439.

Lajoie, S. & Lesgold, A. (1989). Apprenticeship training in the workplace: Computer-coached practice environment as a new form of apprenticeship. Machine-Mediated Learning, 3, 7-28.

Lee, F. & Anderson, J. (2001). Does learning a complex task have to be complex?: A study in learning decomposition. Cognitive Psychology, 42, 267-316.

LeFevre, J. & Dixon, P. (1986). Do written instructions need examples? Cognition and Instruction, 1986, 3, 1-30.

Loftus, G.R. (1993). A picture is worth a thousand p values: On the irrelevance of hypothesis testing in the microcomputer age. Behavior Research Methods, Instruments, & Computers, 25, 250-256.

Logan, G. D. (1991). Automaticity and memory. In W. E. Hockley & S. Lewandowshy (Eds.), Relating theory and data: Essays on human memory in honor of Bennet B. Murdock (pp. 347-366). Hillsdale, NJ: Erlbaum.

Logan, G.D. (1988). Toward an instance theory of automatization. Psychological Review, 95, 492-527.

Mioduser, D., Venezky, R. & Gong, B. (1996). Students' perceptions and designs of simple control systems Computers in Human Behavior, 12, 363-388.

Moley, B., Hart, S, Leal, L., Santulli, K., Johson, T. & Hamilton, L. (1992). The teacher's role in facilitating memory and study strategy development in the elementary school classroom. Child Development, 63, 653-672.

Munley, G., & Patrick, J. (1997). Training and transfer of a structural fault-finding strategy. Ergonomics, 40, 92-109.

Needham, D. & Begg, I. (1991). Problem oriented training promotes spontaneous analogical transfer: Memory orientated training promotes memory for training. Memory and Cognition, 19, 543-557.

Neves, D. & Anderson, J. (1981). Knowledge complilation: Mechanisms for the automatization of cognitive skills. In J. R. Anderson (Ed.), Cognitive skills and their acquisition. Hillsdale, NJ: Erlbaum.

Newell, A. & Rosenbloom, P. (1981). Mechanisms of skill acquisition and the law of practice. In J. R. Anderson (Ed.), Cognitive Skills and their Acquisition (pp. 1-55). Hillsdale, NJ: Erlbaum.

Nickerson, R. (1999). Why interactive computer systems are sometimes not used by people who might benefit from them. International Journal Of Human-Computer Studies,51, 307-321.

Novick, L. (1988). Analogical transfer, problem solving similarity, and expertise. Journal of Experimental Psychology:Learning, Memory and Cognition, 14, 510-520.

O' Reilly, T., & Dixon, P. (2001a). The Role of Functional Labels and Learning Strategy on the Transfer of Prior Device Experience. (Manuscript Submitted to Human Factors).

O' Reilly, T., & Dixon, P. (2001b). The Interface and its Effect on the Transfer of Device Procedures. (Unpublished Manuscript).

Perfetto, G., Bransford, J. & Franks, J. (1983). Constraints on acess in a problem solving context. Memory and Cognition, 18, 83-98.

Pirolli, P, & Anderson, J. (1985). The role of learning form exaples in the acquisitions of recursive programming skills. Canadian Journal of Psychology, 39, 240-272.

Reeves, L. & Weisberg, R. (1994). The role of content and abstract information in Analogical transfer. Psychological Bulletin,115, 381-400.

Reeves, L. & Weisberg, R. (1990, March-April). Analogical transfer in problem solving: Schematic representations and cognitive processes. Paper presented at the meeting of the Eastern Psychological Association, Philadelphia. PA.

Ross, B. (1984). Remindings and their effect on learning a cognitive skill. Cognitive Psychology, 16, 371-416.

Ross, B. (1987). This like that: The use of earlier problems and the separation of similarity effects. Journal of Experimental Psychology: Learning, Memory and Cognition,14, 629-639.

Ross, B. (1989). Distinguishing types of superficial similarities: Different effects on the access and use of earlier problems Journal of Experimental Psychology: Learning, Memory and Cognition, 15, 456-468.

Ross, B. & Kennedy, P. (1990). Generalizing from the use of earlier examples in problem solving. Journal of Experimental Psychology: Learning Memory, and Cognition, 16, 42-55.

Ross, B & Kilbane, M. (1997). Effects of principle explanation and superficial similarity on analogical mapping in problem solving. Journal of Experimental Psychology: Learning Memory, and Cognition, 23, 427-440.

Schroth, M. (1995). Variable delay of feedback procedures and subsequent concept formation transfer. Journal of General Psycholgy,122, 393-399.

Schumacher, R. M., & Gentner, D. (1988). The transfer of training as analogical mapping. IEEE Transactions on Systems, Man, and Cybernetics, 18, 592-600.

Schunn, C. & Dunbar, K. (1996). Priming, analogy, and awareness in complex reasoning. Memory & Cognition,24, 271-284.

Singley, M. & Anderson, J. (1988). A keystroke analysis of learning and transfer in text editing. Human-Computer Interaction, 3(3) 223-274.

Spencer, R. & Weisberg, W. (1986). Context-dependent effects on analogical transfer. Memory & Cognition, 14, 442-449.

Sternberg, S. (1969). The Discovery of Processing Stages: Extensions of Donders' Method. Acta Psychologica, 30, 276-315.

Stein, B., Way, K, Benningfield, S. & Hedgecough, C. (1986). Constraints on spontaneous transfer in problem solving tasks. Memory & Cognition, 14, 432-441.

Thompson, B. (1993). The use of statistical significance tests in research: Bootstrap and other alternatives Journal of Experimental Education, 61, 361-377.

Thorndike, E., & Woodworth, R. (1901a). The influence of improvement in one mental function upon the efficiency of other functions. (I) The Psychological Review, 8, 247-261.

Thorndike, E., & Woodworth, R. (1901b). The influence of improvement in one mental function upon the efficiency of other functions. (II) The Estimation of Magnitudes. The Psychological Review, 8, 384-395.

Thorndike, E., & Woodworth, R. (1901c). The influence of improvement in one mental function upon the efficiency of other functions. (III) Functions Involving Attention, Observation and Discrimination. The Psychological Review, 8, 553-564.

Trudel, C. & Payne, S. (1996). Self-Monitoring during exploration of an interactive device. International Journal of Human-Computer Studies, 45, 723-747.

Van der Veer, G. (1989) Individual differences and the user interface. Ergonomics, 32, 1431-1449.

VanLehn, K. (1996). Cognitive skill acquisition. Annual Review of Psychology,47, 513-539.

VanLehn, Kurt. (1998). Analogy events: How examples are used during problem solving. Cognitive Science, 22, 347-388.

Wharton, C. M., Holyoak, K., Downing, P., Lange, T. Wickens, T. & Melz, E. (1994). Below the surface: Analogical similartity and retrieval competition in reminding. Cognitive Psychology, 26, 64-101.

Weisberg, R., DiCamillo, M. & Phillips (1978). Transferring old associations to new situations: A non automatic process. Journal of verbal Learning and Verbal Behavior, 17, 219-228.

Wertheimer, M, (1959). Productive Thinking. (Enlarged ed.) New York: Harper & Brothers. (Original work published 1945)

Woltz, D., Gardner M. & Bell, B. (2000). Negative transfer errors in sequential cognitive skills: Strong-but-wrong sequence application. Journal of Experimental Psychology: Learning, Memory, & Cognition,26, 601-625.