Training-Transfer Guidelines for Virtual Environments (VE)

LT Joseph Cohn, Ph.D., Jessica Helmick, Christopher Myers
Naval Air Warfare Center Training Systems Division
Orlando, FL 32826

John Burns, Ph.D.
Sonalysts, Inc.
Orlando, FL 32826

ABSTRACT

In recent years, advances in both computer hardware and software have set the stage for designing Virtual Environments (VE) of ever-increasing fidelity. These improvements in VE technology have revived interest in using virtual worlds to provide training. There are many advantages to using VE-based training. For example, VE provides a cost-effective, flexible training environment that can be quickly and easily reconfigured to provide mission-specific training. Also, VE affords instructors the opportunity to expose students to situations that would otherwise be impossible (i.e. life threatening) to recreate in real-life training scenarios. As well, VE provides a unique opportunity for trainers to evaluate their students either in real time, by freezing training at critical points, or by replaying the entire training scenario upon completion.

One of the key assumptions in using VE-based training is that the training received in the VE world will transfer to the real world. However, it has often proven difficult to establish this transfer of training. One reason for this difficulty is that a consensus is lacking in how to establish that training-transfer has occurred. We present here a system of guidelines for establishing training transfer from a VE to a real-world task. In formulating our guidelines we draw upon a wide range of sources, including the flight simulator literature, academic and human factors research as well as findings from our own research.

Issues to be addressed include: defining a specific training task in terms of a series of readily observable variables that are critical to successfully learning the task; providing subjects with training in a VE that emphasizes this variable set; transitioning these trained subjects to the real-world task, while recording these same variables; observing a control group exposed only to the real-world task; finally, using a cross-validation process (Subject Matter Expert feedback), to supplement our evaluation of the degree to which training transfer has occurred. We choose as our model case a shiphandling task. Establishing a set of guidelines should provide future trainers/VE developers with a set of tools for determining how best to design their VE worlds and training protocols.

BIOGRAPHICAL SKETCH

LT Joseph Cohn is an Aerospace Experimental Psychologist at NAWCTSD, currently involved with their Virtual Environment Training Testbed (VETT) program. He received his Ph.D. from Brandeis University's Ashton Graybiel Spatial Orientation Laboratory, where he studied the effects of VE on human motor control. He recently completed a postdoc with Dr. J.A.S. Kelso, at Florida Atlantic University's Center for Complex Systems, where he explored the spatio-temporal aspects of motor learning.

John Burns is a senior scientist at Sonalysts, Inc. working in Orlando, FL. He has been working closely with NAWCTSD in R & D of shipboard training systems, focusing on the development of performance measurement methodologies and technologies. Burns received his Ph.D. in psychology from the University of Massachusetts in Amherst and has 5 years of experience in team training research and development.

Jessica Helmick graduated (Cum Laude) from the University of Central Florida with a BA in Psychology. She is actively involved in NAWCTSD's VETT program, where she is presently working as a Research Psychologist.

Christopher Myers is a senior at the University of Central Florida, pursuing a BS in Psychology. He is actively involved in NAWCTSD's VETT program where he currently works as a Research Assistant.

Training-Transfer Guidelines for Virtual Environments (VE)

LT Joseph Cohn, Ph.D., Jessica A. Helmick, Christopher Myers
Naval Air Warfare Center Training Systems Division
Orlando, FL 32826

John Burns, Ph.D.
Sonalysts, Inc.
Orlando, FL 32826

INTRODUCTION

In 1929, Edwin Link introduced the world to his flight simulator (U.S. Congress, 1994), changing forever the military aviation training landscape. Although his was not the first such trainer (that distinction going to the English, in 1910, with their 'Anders Teacher' and 'Eardly-Billings Oscillator'), it was the first to gain wide-scale use for military training applications (Valverde, 1973). The advent of this new technology raised many questions, foremost of which was how well training in a 'simulation' could substitute for training in the real flight environment. Edwin Link, performing what may have been the first training-transfer experiment in a flight simulator, demonstrated that, after 6 hours in his trainer, his brother required less than one hour of actual flying time to acquire basic flight skills (U.S. Congress, 1994). In the same vein, Link also indirectly performed the first simulator cost effectiveness study, since he was able to train pilots, using his Link Trainer, to solo for the lump sum of $85, while other flight instructors were charging over $200! Nevertheless, after decades of research, the questions of how well training obtained in a simulator transfers to the real world has yet to be decisively answered (Detterman & Sternberg, 1993).

Clearly, one of the greatest benefits of simulated training is the cost-savings that it affords. This benefit is realized in several ways. First, simulators cost less to run than the real devices that they mimic. For example, the Army's M1 Tank Driver Training device has been shown to cost only $6 per simulated mile, versus an estimated $92 per actual mile for using a real M1 (Raisler & Lampton, 1995). Second, simulators can effectively reduce training costs, thereby saving on wear and tear to the actual devices. For instance, practicing a series of actual landings on an aircraft carrier stresses the airframe, thereby shortening its life cycle. Practicing the same landings on a simulator avoids stressing the actual aircraft, keeping it in the loop longer for actual combat (c.f. Westra, et al, 1986). As well, some of the best uses for simulators occur when training under real conditions could be unsafe (Caird, 1995). For example, simulators provide a means for pilots to practice full-scale air- to- air combat maneuvers (Pohlmann & Reed, 1978; Valverde, 1973; Baudhin, 1987) or air- to- ground attack patterns (Lintern, et al, 1984). Simulators also provide trainees with an easily accessible way of practicing portions of an overall mission, without the need for a full work-up. In this way, trainees who feel they are weak in one training area can practice that particular area repeatedly, until they feel secure in their knowledge and skills. In many instances, this could only be done in the real environment at great cost, as in the case of practicing certain aspects of precision aerial maneuvers, like formation flying (Reid & Cyrus, 1974). In addition to providing for the trainees, simulation also affords instructors unique opportunities to train their students in ways that simply can not be matched in the real environment. For example, training in a simulator provides instructors the opportunity to stop the training simulation when they feel the need to point out a particular mistake, or an impending mistake, then review it with the student, and, finally, restart the training at a suitable point (Waag, 1981; Pohlmann & Reed, 1978).

Until recently, most simulators were of the legacy, one-of-a-kind and immobile variety. Vast spaces were necessary to house these expensive-to-operate, platform-specific training devices, and they required a host of maintenance personnel to assist in their operation. The introduction of Virtual Environments (VE) marked a watershed in simulation technology. In addition to the cost-savings inherent in any type of simulation, VEs have the added benefit of being more portable than traditional training devices. In fact, current hardware technology allows us to package a tremendous amount of computing potential in small desktop systems. As a case in point, we have developed a Virtual Environment Training Testbed (VETT), consisting primarily of 2 Silicon Graphics CPUs, 2 Monitors, a Head Mounted Display and a microphone (Figure 1). This device requires only two engineers, working for one hour, to set up at virtually any location. Thus, VE simulators can be brought to the trainees, rather than bringing the trainees to the simulation. Additionally, the ever-decreasing cost of desktop computer hardware essentially guarantees that VE simulations will cost less to construct, modify and run than legacy type trainers.

Figure 1: Photo of equipment for running VETT. Note HMD with head tracking unit on top and microphone on the subject, with monitors in the background. Not shown: 2 CPUs located beneath the monitors.

Another benefit of VE simulators is that they are software driven. This means that changes to the simulation require simple modifications to the existing code. Consequently, one VE training station could service a range of training needs, for multiple platforms. For instance, once the scenery is mapped onto the virtual world, different modules representing different types of watercraft, such as submarines, oilers and landing craft can all be interchanged in this environment. The only pieces of the VR puzzle that would need to be stipulated would be the image mapping of the specific craft together with physical models describing how each craft should perform in the virtual world (hydrodynamic or aerodynamic models). In addition, many of the ‘control devices’ that one would expect to find in each of these craft (dials indicating speed and so forth) can actually be rendered virtually. In contrast, traditional trainers require a significant amount of physical equipment (such as mock-ups of bridges, or even whole cockpits complete with buttons, switches and knobs and dials), most of which would need to be changed for each type of craft being simulated.

THE ISSUE

The Need to Re-Examine Training Transfer

Perhaps the most critical assumption underlying the use of simulations is that the training received in these devices transfers effectively to the real world devices that they simulate (Valverde, 1973). The classic design for demonstrating transfer is one in which trainees in an experimental group first receive training in a simulated task, and then transition directly to the real task, while trainees in a control group either perform some other, non specific task prior to the real task, or go straight to the real task (Murdock, 1957; Baudhuin, 1987; Boldovici, 1987). Performance between the two groups is then compared. If the experimental group shows significant performance differences (in a positive sense) over those in the control group, then transfer is said to have occurred.

Although VE simulations have been in use for several decades, it is only within recent years that the technology has advanced to the point that they can begin to provide an immerssive experience. Consequently, the training that students obtain in these advanced devices differs fundamentally from that received in either earlier-type VEs as well as in legacy-type trainers. As a result, with the renewed interest of the military in integrating these more immerssive VE simulations into a range of training domains, it is imperative that we re-examine the issue of training transfer.

A New Method of Analysis

Along with the necessity for this re-examination comes the opportunity for developing a new method of analyzing/demonstrating training transfer. One of the greatest barriers to showing training transfer arises in determining which metrics to use and how best to analyze them. Often, such variables are posed in terms of how many fewer hours the simulation-exposed trainees required in the actual device, in relation to those denied this training, in order to achieve some performance criterion. For example, in flight simulator evaluations, transfer effectiveness is often described in terms of the ‘aircraft hours savings’, that is, how many less hours in a real aircraft trainees require following exposure to flight simulation training (Reid & Cyrus, 1974; Waag, 1981). In a similar vein, another commonly used variable is the percent savings in terms of trials needed to reach some predetermined level of proficiency (Johnson & Stewart, 1999; Martin & Waag, 1978; Weitz & Adler, 1973; Witmer, et al 1996; Witmer & Sadowski, 1998). Other variables include those derived from subjective evaluation by Subject Matter Experts (SME) (Bell & Waag, 1998; Pohlmann & Reed, 1978) or through the use of questionnaires, administered to the trainees to gain insight into how effective the training is (Pohlmann & Reed, 1978). A quick review of these variables suggests that they are plagued by one of two problems. Either they are entirely subjective, as is the case for the SME and questionnaire data, or they fail to give a full indication of the evolution of the training transfer.

The difficulty with these subjective variables is that they often lack the generalizability that is so crucial for psychology. For example, consider a task that will be expanded upon shortly, Underway Replenishment (UNREP) in which a ship handler is required to come alongside an Oiler, to obtain supplies. Each ship's skipper has their unique way of executing this maneuver. Consequently, a skipper from one ship might rate a Surface Warfare Officer's (SWO) performance as low, while another skipper may rate his/her performance as high. As a result, when we analyze these scores, it is difficult, if not impossible, to extract generalizable principles from them that can then be applied to help us determine how best to enhance training transfer.

As well, many of these variables do not provide researchers and instructors with the full picture of how the training regimen effects changes in performance. Researchers in fields as diverse as motor control, cognitive psychology and perceptual psychology have begun to realize that some of the most important insights into human behavior are obtained by using the relatively new method of time series analysis to investigate how certain behavioral measures change over time (Kelso, 1995; Kelso et al, 1988; Strogatz, 1994). In this way, it is possible to perform a wide range of analyses, including determining the time course behavior of individual variables as well as the temporal relationships between variables. This, in turn, makes it possible to determine more accurately when (and if) certain training interventions should be applied, and for how long. Current analyses on training transfer often focus only on 'lump' sum variables, such as how many wrong turns were made (Witmer & Sadowski, 1998), or how much time was saved using a particular training simulation, instead of actual time in-craft (Reid & Cyrus, 1974). While these measures are important, they fail to illuminate the underlying aspects of training that are transferring, as well as those that aren't. Yet, armed with this deeper knowledge, we might be able to produce increasingly sophisticated training simulations, and to incorporate them into more meaningful training paradigms that will lead to improved transfer of training.

An example of how a time series-based inspection of the data can enhance training is illustrated by an experiment in which subjects were asked to rhythmically move their index finger in the horizontal plane, with the added requirement that peak flexion occur in synchrony with a metronome beat (Cohn and Kelso, 1999). A simple analysis of the data in terms of the time course of finger position suggested that the amplitude of these movements varied inversely with the metronome frequency. A more advanced analysis, in which task performance, defined as the temporal relationship between finger movement and metronome beat (the relative phase variable) was studied, indicated that at both low and high metronome frequencies, subjects were unable to accurately synchronize peak flexion with the metronome. Moreover, correlations between the relative phase variable and a range of other variables indicated that successful performance of the task depended on when, during the course of the movement, the nervous system began decelerating the finger. Note that simply looking at measures such as 'how long it takes to learn a task' or 'how often the subjects performed the task accurately' would not have revealed this information.

THE EXPERIMENT

Guidelines for Determining Training Transfer

Taking our cue from this novel research approach, we have established basic, generalizable guidelines for a method of data analysis that should help establish whether or not training transfer has occurred. The first point is to select a task which is essential to military training, which can benefit from the development of immerssive VE and that provides us with a potentially large number of objective variables that can be easily recorded as they evolve over time, and then be subsequently analyzed. Next, the appropriate training transfer design must be implemented. Following this, the data should be analyzed at the descriptive level. This includes basic plots of the time series of each variable, an indication of task performance and so forth. Then, a more advanced analysis must be performed, in which the temporal relationships between the variables are inspected, and then analyzed in terms of the relative phase, the variability of the relative phase, and possible correlations between this measurement and other variables. In this way, investigators will not only get a deeper intuition for how subjects' performance changes over time, but they will also gain a better understanding of how changes in variables impact each other.

Example Design

One task that is admirably suited to this type of analysis is the UNREP. UNREPs are performed at sea, when one ship (Ownship) requires supplies, equipment and so forth, and receives it from another ship, most usually an Oiler. A typical UNREP involves several phases: the *Approach* phase, in which Ownship begins to move in towards the Oiler; the *Alongside* phase, in which Ownship matches and maintains the course and speed of the Oiler, connects to the Oiler via lines (Phone and Distance lines) and maintain a constant lateral separation; the *Pre Breakaway* and *Breakaway* phases, in which the connecting lines are returned to Ownship and Ownship moves away from the Oiler in a pre-determined fashion (Figure 2). The SWO stands on the bridge monitoring Ownship's progress, and effects the necessary changes by issuing verbal commands, such as 'hard left' to steer Ownship and 'increase turns' to change Ownship's speed. The task is, therefore, primarily perceptual in nature, with a verbal output component. Smooth and timely accomplishment of this task is of paramount concern to a SWO. Failure in any one of these maneuvers could result in, at the very the least, a delay in obtaining essential supplies, and at worst could mean damage to multi-million dollar vessels and even loss of life. In the past, SWOs primarily learned this task on-the-job, a method which is both costly in terms of using fleet vessels and leaves little room for 'beginner's error'.

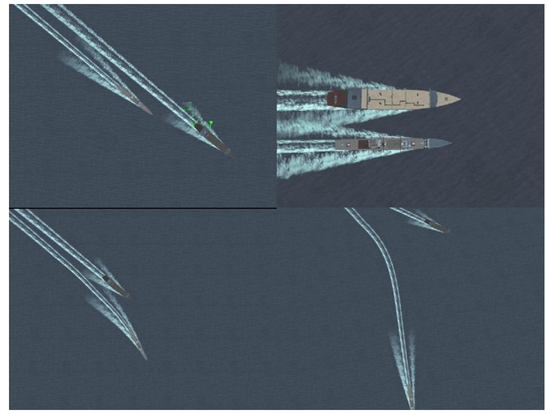

Figure 2. 4 Phases of the UNREP. Top Left to Right: *Approach, Alongside*. Bottom, Left to Right: *Pre-Breakaway, Breakaway*.

The UNREP task allows us to collect and analyze a number of meaningful, objective variables during training as well as during performance of the actual task, which can then be analyzed for improvement during training as well as for transfer to the real world. These variables include: Range of Ownship from the Oiler; Course of Ownship; Speed of Ownship; and the types of commands issued.

For our purposes, a suitable design would require two groups of newly designated SWOs. One group, the control group, will receive UNREP training directly in the real environment. A second group, the treatment group, will first practice a series of UNREP maneuvers using the VETT training system, and will then perform UNREPs in the real environment.

At the conclusion of the experiment we will analyze the data using the two approaches mentioned previously. At the descriptive level, the recorded variables can simply be plotted out, and inspected in terms of how often the trainee was able to maintain required values of these variables and for how long. At a deeper level, we can examine the temporal relationship between either Range, Course or Speed and the commands issued. Initially, we might expect this relationship to be unstable, with a large amount of variability, indicating that subjects are unsure of when to time their commands in order to effect the desired change. With increased practice, however, we expect this relationship to become more stable.

Current Efforts

It was our intent to first adapt the time series analysis methods using data obtained from a single SWO, run through multiple training trials. Once the methods had been sufficiently established, we could then analyze data from a full-scale experiment. Thus, as a first pass at analyzing data in this manner, we collected training data from a single SWO performing 10 UNREP trials, using the VETT system at the Naval Air Warfare Center, Training Systems Division. The simulated Ownship was a Spruance-class Destroyer (DD-963). Trial 2 was interrupted mid-run, but is presented for the sake of continuity. Data were collected for five days over a two-week period, with the participant performing two trials per day. Each trial lasted approximately 600 seconds. The data were collected in real time, filtered and subsequently output at 1Hz. Although the participant did not go on to perform the actual UNREP at sea, future experimental design would ultimately require this in order to show transfer.

RESULTS AND ANALYSIS

For each trial, Ownship was placed approximately 2100 feet (700 yards) behind the Oiler, requiring the participant to first bring Ownship within the vicinity of the Oiler, as indicated by the steep initial slope present in each plot (Figure 3). The range goal was to maintain an intership distance of 120 feet. The important point to notice in scanning across these plots is the change that occurred with training. Note that in the early trials (1-4) the participant appears to drive Ownship to within approximately 150 feet, but then drives the ship far out of range by the end of the trial. On the other hand, by the end of the 10th trial, the participant appears to maintain a steady, albeit slightly outside the desired parameters, range throughout. A similar trend is found with the participant's ability to maintain the required course heading of 130 degrees (Figure 4). In the earlier trials, the participant appears to 'bounce around' throughout the trial, only occasionally settling on the required heading. By later trials, although the participant still was unable to maintain a constant heading of 130 degrees, his excursions from this value are far less pronounced. Finally, the participant had difficulty maintaining the required 15 knot speed in the earlier trials, but, with additional trials, was better able to do so (Figure 5).

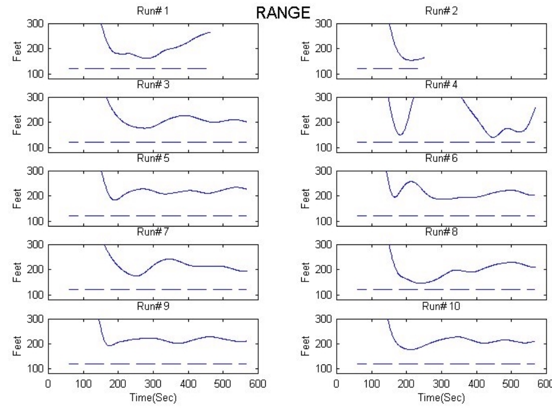

Figure 3. Range time series. Dashed line indicates required range value (120 feet).

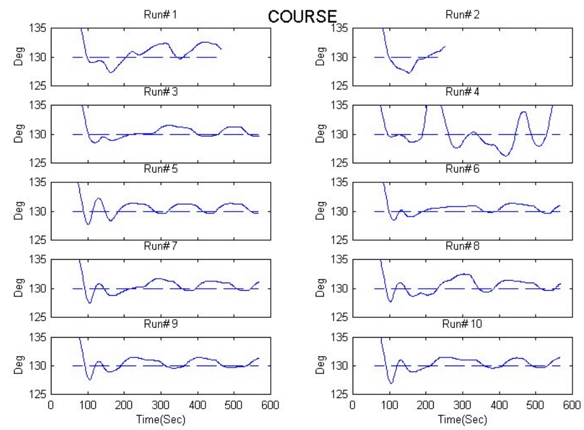

Figure 4. Course time series. Dashed line indicates required course value (130 degrees).

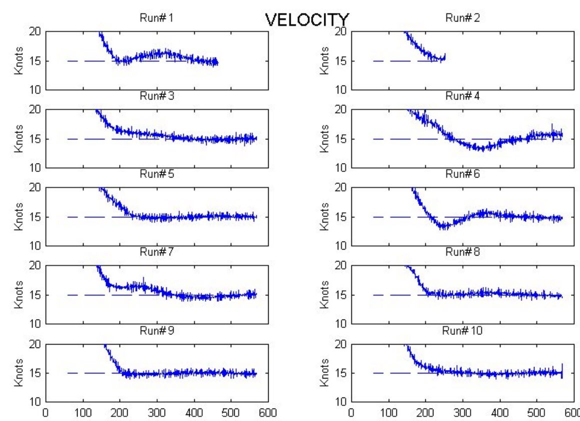

Figure 5. Speed time series. Dashed line indicates required speed value (15 knots).

A complementary analysis looks at how often the participant was within acceptable limits for range, course and speed, and how long, on average, the participant stayed within these limits. Acceptable limits were defined as allowable deviations from the ideal value, based on an inspection of the data for each of these three variables. These values, for the current analysis, were: range, +/- 36 feet; course, +/- 1 degree and speed +/- 1.5 knots. Although the participant was able to achieve these values frequently, for the most part he held them for only several seconds. We therefore chose to exclude those events that lasted less than a criterion amount of time, as determined by inspecting the data. These values were: range, 15 seconds; course, 15 seconds; speed, 25 seconds.

The participant was within acceptable range limits for three of the ten trials (2, 4 and 8), and this occurred once per trial (Figure 6, bottom). Since there was only one such event per trial, the average time is simply the duration of this event. Interestingly, the time spent within this acceptable range increased as training progressed, from 24 seconds on trial 2 to 64 seconds on trial 8 (Figure 6, top). The average time spent within acceptable course limits fluctuates from about 65 seconds on the first trial, to about 136 seconds on the 10th trial (Figure 7, top). The participant was within course limits at least two times, for every trial (Figure 7, bottom). This evaluates to 65 sec*3=195 seconds total time spent within the correct limits, on trial 1, and 136*3=408 seconds on the last trial, indicating again that participant's performance was enhanced through repeated exposure to this task. Although the number of times the participant was within speed limits decreased, the average amount of time spent per event increased. Thus, on trial 1, the participant was within limits twice, for 95 seconds (=190 seconds total), while on trial 10 the participant was only within limits once, but for almost 400 seconds (essentially, 67% of the entire trial!), demonstrating marked improvement (Figure 8 bottom, top).

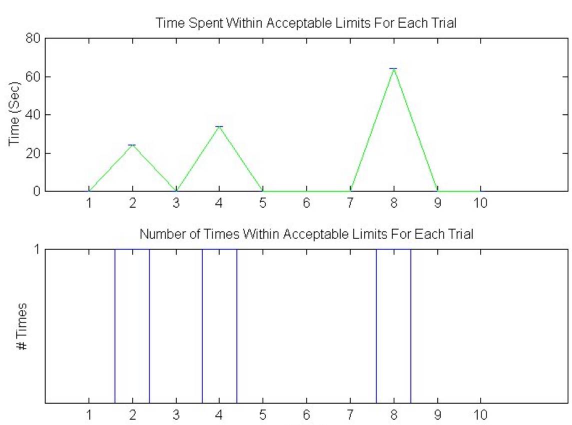

Figure 6. Analysis for range.

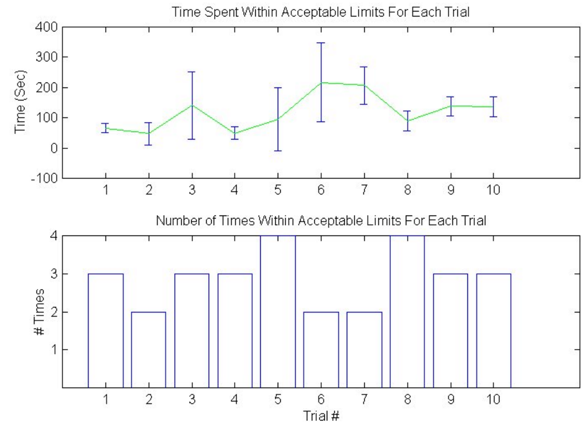

Figure 7. Analysis for course.

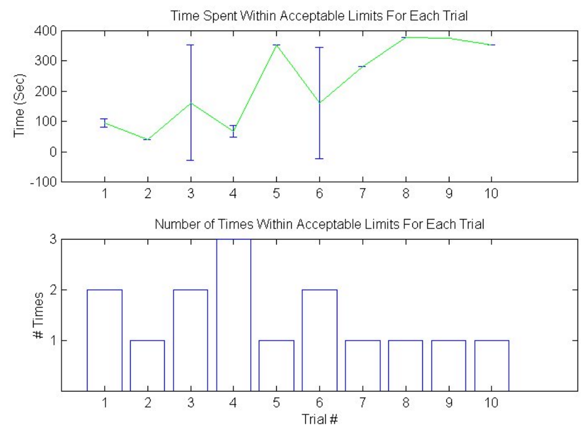

Figure 8. Analysis for speed.

Since the UNREP is primarily a perceptual/verbal task, where the SWO drives the ship by issuing orders, an additional measurement of the participant's performance is obtained through an analysis of the voice commands issued during the maneuver. These commands are broken down into four main categories: Informational (e.g.: requesting range, course and speed values), Engine (change speed), Course (change course) and Other (commands that don't fall into these three areas). The amount, type and relevance of the commands issued serves as a good indicator of whether or not the SWO is learning to perform the task well. A SWO who issues a number of Engine commands in order to match the Oiler's speed is considered to be a much poorer performer than a SWO who issues only a few such commands to achieve the same goal. As well, an accomplished SWO is able to issue Engine and Course commands while requesting only a minimal amount of information, being able to judge the necessary values perceptually.

The data are analyzed in terms of the frequency with which each type of command was issued. Looking for trends across trials, we can gain better insight into the time evolution of the training (Figure 9). Note that the total number of commands, posted above each plot, decreases with successive trials. This suggests that the participant was able to achieve the necessary range, course and speed with successively less commands, indicating improvement with training. Additionally, for all trials, the participant was able to issue Engine and Course commands without resorting to informational requests, suggesting that the participant relied mainly on perceptual cues to perform the UNREP.

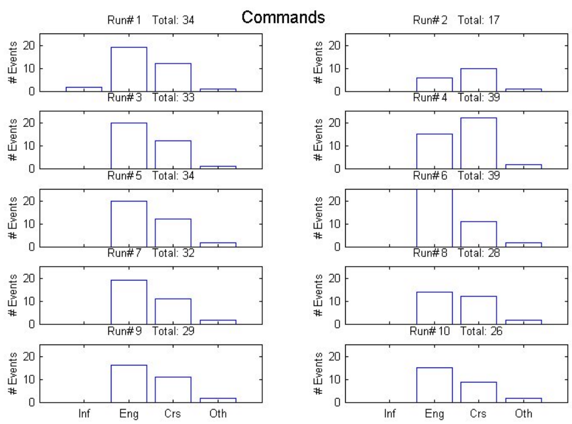

Figure 9. Histogram analysis of four different types of commands issued (Inf: Informational; Eng: Engine; Crs: Course; Oth: Other).

CONCLUSION

Performance On The UNREP

Taken together, our preliminary analysis suggests that with only a limited number of trials, the participant's performance on the simulated UNREP task improved, as determined by analyzing the range, course and speed variables. By the 10th trial, the participant was able to stay within the required range for over 60 seconds, and to maintain the required course and speed for nearly 400 seconds. Moreover, the observed improvements were coincident with the participant issuing commands a fewer number of times. Most importantly, by looking at the time series plots for a set of variables, we determined that the participant initially had a tendency to pull away from the Oiler (range), had trouble maintaining a constant, correct heading (course), and tended to either overshoot or to undershoot the speed (speed). Armed with this information, we can better evaluate when, during each trial, feedback might aid the participant. Moreover, our proposed method of analysis also lets us better estimate just what the content of the feedback should be. Specifically, rather than simply relying on a SME score to tell us how well the participant did, we can now provide the participant with exact detail, such as 'you have a tendency to pull away from the Oiler towards the end of the trial' or 'your course fluctuates wildly at the start of the trial'. Providing information at this level should help trainees learn the task more completely, at a faster rate. Additionally, our method of analysis allows us to more accurately determine when training should be considered complete. Looking at figures 3 through 5, it is clear that 10 trials are not sufficient to bring the participant within acceptable range limits, but may be enough to bring him within acceptable course and speed limits. The trainer can then make the decision to terminate the training or to provide additional training, emphasizing the range aspect. Finally, once the time series for these variables are in line with successfully completing training, the participant can then perform the UNREP using real ships. In order to demonstrate the transfer of the VE training to the real task, the same variables would be measured and analyzed, as in the training portion. Measurements indicating the degree of similarity between the very first real world UNREP trial and the very last VE UNREP trial would be used to indicate the degree to which training has transferred to the real world. These findings could then be enhanced by more traditional performance measures, such as evaluation by SMEs.

Applicability to Other Training Domains

The analytical guidelines discussed here are not unique to the UNREP task. Rather, there are various VE training domains whose potential for transfer to the real world can be resolved by these methods. In order to do so, the investigator must properly identify a key task, the relevant variables that describe it and the expected relationships between these variables. For instance, consider an aircraft carrier landing trainer. A key task may involve the final stages of recovery, in which the pilot must work his plane into the 'groove'. To stay within this narrow region, the pilot must make fine adjustments to the plane's power and attitude. This, then, is our task. Candidate variables include: position relative to ship; angle of attack and frequency/magnitude of power adjustments. Using the method outlined above, we might first simply look at the time series for each of these variables, looking for points at which sharp changes to any of these variables occurred. Next, we may look at the temporal relation (relative phase) between angle of attack and power adjustments, or even position relative to ship and power adjustments. These variables might tell us whether, over time, the subject was able to learn how best to jockey power in order to stay in the 'groove'. Ultimately, the information obtained through the application of these analysis guidelines should provide investigators with an additional tool for improving training regimens and for determining whether or not a given VE simulation provides for transfer to the real world environment.

REFERENCES

Baudhuin, E. S. (1987). The design of industrial and flight simulators. In S. M. Cormier (ed.), *Transfer of learning: contemporary research and learning* (San Diego: Academic Press, Inc.) pp.217-237.

Bell, H. H. and Waag, W. L. (1998). Evaluating the effectiveness of flight simulators for training combat skills: a review. *The International Journal of Aviation Psychology*, 8(3), 223-242.

Boldovici, J. A. (1987). Measuring transfer in military settings. In S. M. Cormier (ed.), *Transfer of learning: contemporary research and learning* (San Diego: Academic Press, Inc.) pp.239-260.

Caird, J. K. (1995). *Issues in the application of virtual environment systems to military training and selection*. (Tech. Rep. 9-94). Toronto: Canadian Forces, Personnel Applied Research Unit.

Cohn, J. V and Kelso, J.A.S. (1999). On getting your finger to the right place at the right time. (Paper presented at *1999 North American Society for Psychology of Sport and Physical Activity conference*).

Detterman, D. K. and Sternberg R.J. (1993). *Transfer on trial: intelligence, cognition, and instruction* (Norwood, NJ: Ablex Publishing Corporation).

Johnson, D. M. and Stewart II, J. E. (1999). Use of virtual environments for the acquisition of spatial knowledge: comparison among different visual displays. *Military Psychology*, 11(2), 129-148.

Kelso, J. A. S. (1995). *Dynamic patterns: the self-organization of brain and behavior* (Cambridge, MA: The MIT Press).

Kelso, J. A. S., and Jeka, J. J. (1992). Symmetry breaking dynamics of human multilimb coordination. *Journal of Experimental Psychology Human Perception and Performance*, 18, 645-668.

Kelso, J. A. S., Scholtz, J. P., & Schoner, G. (1988). Dynamics governs switching among patterns of coordination in biological movement. *Physics Letter A, 134, 8-12.*

Lintern, G., Thomley, K. E., Nelson, B. E., & Roscoe, S. N. (1984, July). *Content, variety, and augmentation of simulated visual scenes for teaching air-to-ground attack* (NAVTRAEQUIPCEN 81-C-0105-3). Naval Air Equipment Center, FL: Defense Technical Information Center. (AD A145 218).

Martin, E. L. and Waag W. L. (1978). *Contributions of platform motion to simulator training effectiveness: study1-basic contact* (AFHRL-TR-78-15). Brooks Air Force Base, TX: Air Force Human Resource Laboratory.

Murdock, B. B. (1957). Transfer design and formulas. *Psychological Bulletin*, 54(4), 312-326.

Pohlmann, L. D. and Reed, J. C. (1978, October). *Air-to-air combat skills: contribution of platform motion to initial training* (AFHRL-TR-78-53). Brooks Air Force Base, TX: Air Force Human Resource Laboratory.

Raisler, R. B. and Lampton D. R. (1995, October). *Simulator sickness in tank driver trainers*. Army Institute for the Behavioral and Social Sciences, VA. (AD A304277).

Reid, G. B. and Cyrus, M. L. (1974, December). *Transfer of training with formation flight trainer* (AFHRL-TR-74-102). Brooks Air Force Base, TX: Air Force Human Resource Laboratory. (AD A009638).

Strogatz, S.H. (1994). *Nonlinear dynamics and chaos: with applications to physics, biology, chemistry and engineering.* (Reading, MA: Addison-Wesley).

U.S. Congress, Office of Technology Assessment, *Virtual Reality and Technologies for Combat Simulation—Background Paper*, OTA-BP-ISS-136

(Washington, DC: U.S. Government Printing Office, September 1994).

Valverde, H. V. (1973). A review of flight simulator transfer of training studies. *Human Factors,* 15(6), 510-523.

Waag, W. L. (1981, January). *Training effectiveness of visual and motion simulation* (AFHRL-TR79-72). Brooks Air Force base, TX: Air Force Human Resource Laboratory. (No. AD A094530).

Weitz, J. and Adler, S. (1973). The optimal use of simulation. *Journal of Applied Psychology*, 58(2), 219-224.

Westra, D. P., Lintern, G., Sheppard, D. J., Thomley, K. E., Mauk, R., Wightman, D. C., & Chambers, W. S. (1986, June). *Simulator design and instructional features for carrier landing: a field transfer study* (NAVTRASYSCEN 85-C-0044-2). Naval Training Systems Center, FL: Naval Air Systems Command. (No. AD A169962).

Witmer, B.G., Bailey, J.H., & Knerr, B. W. (1996). Virtual spaces and real world places: transfer of route knowledge. *International Journal Human-Computer Studies,* 45, 413-428.

Witmer, B.G., and Sadowski, Jr., W. J. (1998). Nonvisually guided locomotion to a previously viewed target in real and virtual environments. *Human Factors*, 40(3), 478-488.