

# **SANDIA REPORT**

SAND2014-4762  
Unlimited Release  
Printed June 2014

## **The Effect of Job Performance Aids on Quality Assurance**

Erik Fosshage

Prepared by  
Sandia National Laboratories  
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.



**Sandia National Laboratories**

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

**NOTICE:** This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from

U.S. Department of Energy  
Office of Scientific and Technical Information  
P.O. Box 62  
Oak Ridge, TN 37831

Telephone: (865) 576-8401  
Facsimile: (865) 576-5728  
E-Mail: [reports@adonis.osti.gov](mailto:reports@adonis.osti.gov)  
Online ordering: <http://www.osti.gov/bridge>

Available to the public from

U.S. Department of Commerce  
National Technical Information Service  
5285 Port Royal Rd.  
Springfield, VA 22161

Telephone: (800) 553-6847  
Facsimile: (703) 605-6900  
E-Mail: [orders@ntis.fedworld.gov](mailto:orders@ntis.fedworld.gov)  
Online order: <http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online>



SAND2014-4762  
Unlimited Release  
Printed June 2014

# The Effect of Job Performance Aids on Quality Assurance

Erik Fosshage  
System Surety Engineering  
Sandia National Laboratories  
P.O. Box 5800  
Albuquerque, New Mexico 87185-MS0638

## Abstract

Job performance aids (JPAs) have been studied for many decades in a variety of disciplines and for many different types of tasks, yet this is the first known research experiment using JPAs in a quality assurance (QA) context. The objective of this thesis was to assess whether a JPA has an effect on the performance of a QA observer performing the concurrent dual verification technique for a basic assembly task. The JPA used in this study was a simple checklist, and the design borrows heavily from prior research on task analysis and other human factors principles. The assembly task and QA construct of concurrent dual verification are consistent with those of a high consequence manufacturing environment. Results showed that the JPA had only a limited effect on QA performance in the context of this experiment. However, there were three important and unexpected findings that may draw interest from a variety of practitioners. First, a novel testing methodology sensitive enough to measure the effects of a JPA on performance was created. Second, the discovery that there are different probabilities of detection for different types of error in a QA context may be the most far-reaching results. Third, these results highlight the limitations of concurrent dual verification as a control against defects. It is hoped that both the methodology and results of this study are an effective baseline from which to launch future research activities.

## **ACKNOWLEDGMENTS**

This master's thesis is a collaboration between the Surety Assessment, Engineering, and Analysis Center (0400) at Sandia National Laboratories (SNL) and the School of Industrial Engineering at Purdue University. Funding was provided by the University Part Time program at SNL as well as the Nuclear Safety Research and Development Program within the National Nuclear Security Administration of the United States Department of Energy.

The author wishes to thank Professor Barrett Caldwell at Purdue University and Caren Wenner at Sandia National Laboratories for technical guidance and direction during this research. Additional advice at SNL was provided by Courtney Dornburg, Susan Stevens-Adams, Kerstan Cole, Judi See, Ed Thomas, Jeffrey West, Jennifer Samora, Jenny Lindblom, and Leslie Simmons. Administrative support was provided by Shelby Hopkins, Gwen Zon, Lori Hatcher, Julie Hillskemper, Laura Calderone, Missy Hess, Susan Fourt, Tiffany Vargas, Amy Rein, Rogulja Wolf, Gary Boulanger, Tom Wubbels, and Ann Osburn.

Lastly, the author has been honored to engage in a correspondence with Dr. Edgar L. Shriver; a true pioneer in the human factors community regarding task analysis and job performance aids.

## CONTENTS

1. INTRODUCTION .....	7
1.1 Motivation.....	7
1.2 Research Statement.....	8
1.3 Thesis Organization .....	8
2. LITERATURE REVIEW .....	9
3. METHODOLOGY .....	13
3.1 Experimental Task Selection .....	13
3.2 Test Approach.....	15
3.3 Sample Size Determination.....	17
3.4 Test Procedure .....	18
3.5 Job Performance Aid.....	22
4. RESULTS AND ANALYSIS.....	25
5. DISCUSSION.....	31
6. CONCLUSIONS AND FUTURE RESEARCH .....	33
7. REFERENCES .....	35
APPENDIX A. JOB PERFORMANCE AID .....	41
APPENDIX B. FAULT TYPES.....	43

## FIGURES

Figure 1. Test Setup .....	19
Figure 2. Sample Instructions Page, Depicting Fault #1 (Fault Type 1) .....	21
Figure 3. Fault #1 Occurring in an Experimental Trial .....	22
Figure B1: Sample Instructions Page, Depicting Fault #1 (Fault Type 1).....	43
Figure B2: Fault #1 Occurring in an Experimental Trial.....	44
Figure B3: Sample Instructions Page, Depicting Fault #8 (Fault Type 2).....	45
Figure B4: Fault #8 Occurring in an Experimental Trial.....	46
Figure B5: Sample Instructions Page, Depicting Fault #3 (Fault Type 3).....	47
Figure B6: Fault #3 Occurring in an Experimental Trial.....	48
Figure B7: Sample Instructions Page, Depicting Fault #5 (Fault Type 4).....	49
Figure B8: Fault #5 Occurring in an Experimental Trial.....	50

## TABLES

Table 1: Definitions Used in the Assembly Task .....	15
Table 2: Probability of Correctly Concluding that $p_2 > p_1$ .....	18
Table 3: Grouping of Experimental Participants By Sequence of Assembly .....	19
Table 4. Fault Descriptions/Types .....	20
Table 5: Fault Detection Performance, By Participant .....	25
Table 6: Performance, By Fault Number .....	26
Table 7: Logistic Regression Table, Pattern A .....	27
Table 8: Logistic Regression Table, Pattern B .....	27
Table 9: Estimated Probability of Detection, Pattern A .....	28
Table 10: Estimated Probability of Detection, Pattern B.....	29

# 1. INTRODUCTION

According to James Cantrell, the main engineer for the Skipper satellite, “It’s always the simple stuff that kills you.” Skipper failed one day into its mission because Russian scientists mistakenly connected the solar panels backwards (“Russians Miswire Satellite,” 1996), an oversight that could have been avoided with the use of a quality assurance (QA) checklist. A checklist may also have prevented the NASA Genesis probe from crash-landing in 2004 when its drogue parachute failed to deploy due to a deceleration sensor that was installed backwards (“Official: Genesis Pre-Launch,” 2006). These examples remind us that even simple human errors can have large consequences, and though human error can never completely be eliminated we can nonetheless introduce processes and tools that are designed to prevent errors or to catch errors and minimize their impact.

One purpose of checklists is to support either simple or complex actions that are subject to the limitations of human performance. In the delivery of a product or activity, an overarching goal for quality assurance and human factors professionals is often the prevention of defects, and this becomes equally as important as it is difficult when dealing with both complex and critical systems. People who are present during high consequence failures (or near misses) always remember the story and are extremely unlikely to repeat the same mistakes, a concept best captured by Vilfredo Pareto in 1870: “Give me a fruitful error any time, full of seeds, bursting with its own corrections” (from *Comment on Kepler*, quoted in Helmreich and Merritt, 2001). A checklist often is a lesson intended for an audience that was not present during a mishap who should learn the failure mechanisms in order to avoid them. This study examines how a checklist (a specific type of job performance aid, or JPA) may affect human performance in the verification activities that are often employed in systems that have high consequences for failure.

## 1.1 Motivation

The author’s original intent in selecting a thesis topic was to bridge the QA and human factors disciplines, because his work experience resides in the former and major field of study is in the latter. Consultations with staff in the Human Factors department at Sandia National Laboratories (SNL) revealed an internal paper titled “The Effectiveness of Verification in Industrial Applications” (Stevens-Adams et al., 2013) that reviews the academic and applied literature on the verification construct, variables underlying its effectiveness, and the applications that are best suited to a particular work environment. The paper cites different definitions of verification from Department of Energy (DOE) standards and handbooks, including the concurrent dual verification technique that is utilized by the Surety Engineering Group<sup>1</sup> at SNL in its high consequence manufacturing environments (DOE, 1993; DOE, 2009) whereby a QA observer is present to witness activities being performed. In consultation with SNL and Purdue advisors, it was determined that there has been a dearth of research on the concurrent dual verification technique and that an appropriate thesis topic would be to address this particular QA task.

---

<sup>1</sup> “Surety” encompasses a number of disciplines related to nuclear weapons such as safety, reliability, quality, human factors, and assessment, and is a term not widely used outside of SNL. The Surety Engineering Group is staffed primarily by quality engineers.

An episode from the author's past involved mentoring an inexperienced QA checker to verify assembly tasks for a high consequence product, including questions about how to specifically complete these tasks. This particular incident related to the performance of concurrent dual verification led to the selection of a job performance aid as the focus of this thesis, and more specifically an examination of checklist aids in QA. After performing a cursory literature search on JPAs, it was initially determined that the topic of JPAs related to QA tasks had not previously been studied. A more thorough search of the JPA literature eventually located a detailed experiment (Shriver et al., 1982) that studied the effect of JPAs on eight (8) different tasks within a nuclear power plant, one of which was specific to QA. However, the QA-focused JPA was eventually dropped from the Shriver study as being too task-interdependent and reliant upon wider functional understanding from an expert operator. The project director (L. Zach, personal communication, June 2013) and lead principal investigator (E.L. Shriver, personal communication, June 2013) of the 1982 report have been located and verified that, to their knowledge, no other research on JPAs for QA has ever been published.

This research study is therefore a nexus for multiple interests: the design and use of checklists to improve human performance, the linkage of human factors and QA techniques for defect prevention, and a continuation of both previous and existing research. In addition to academic discoveries, this study of JPAs can benefit a wide-ranging audience of practitioners as a tool for reducing the likelihood of high consequence events that still happen in every-day life; rockets and satellites continue to fail (Cheng et al., 2009), medical errors lead to tragic results (Dhillon, 2008), and in fact lives could even be saved by implementing a checklist (Haynes et al., 2009). This topic has such an impact that it has even spawned a New York Times bestseller *The Checklist Manifesto* (Gawande, 2010) that has drawn popular interest on an international level.

## **1.2 Research Statement**

Since this was the first known study that examined the use of JPAs for the QA role, the primary research focus was to determine whether a simple JPA (i.e., checklist) has an effect on QA performance, and then only on the concurrent dual verification technique. Once the experimental design was chosen, additional hypotheses were tested regarding independent variables and an important test assumption that was inadvertently germane to the results.

## **1.3 Thesis Organization**

This thesis consists of six chapters. Chapter 2 examines prior research on JPAs in a variety of disciplines, especially human factors, and provides definitions of concurrent dual verification in a QA context. Chapter 3 explains the methodologies used in the design of the experiment, as well as the JPA itself. Chapter 4 presents results and analysis of the data while Chapter 5 discusses the key implications of the study. Chapter 6 presents conclusions and recommends future studies regarding the effect of different JPA formats on QA performance, as well as their potential to affect other QA tools and processes besides verification techniques.

## 2. LITERATURE REVIEW

Probably the best-known JPA is the pilot's checklist, which was born on October 30, 1935 at Wright Air Field in Dayton, Ohio during the initial test flight of the Boeing Model 299 aircraft (which later became the B-17 "Flying Fortress"). The Model 299 crashed after take-off, killing two of the 5 crewmen. The accident investigation board determined that no mechanical failures had occurred and that the aircraft crashed because the pilot forgot to unlock the elevator and rudder locks. The Model 299 was substantially more complex than previous aircraft and simply too much for one pilot to remember how to fly, so the approach at the time was to simplify this complexity with a pilot's checklist (Meilinger, 2004; Gawande, 2010).

Over the next few decades research on JPAs would be conducted in a variety of disciplines, beginning with U.S. Air Force behavioral research in the 1950s. From 1951 to 1956 the Air Force Behavioral Science Laboratory, at the re-named Wright-Patterson Air Force Base in Dayton, contracted the human factors specialists at the American Institute for Research to study how behavioral and psychological factors could be identified and then used to design training requirements for specific tasks (Miller, 1953). It was during this study that Dr. Edgar Shriver coined the term "task analysis" to specify the need to identify stimulus elements, or behavioral cues, that indicate when a task is complete and then place these directly into training instructions (Shriver et al., 1982). Newman (1957) and Miller (1956), under contract to the Air Force Personnel and Training Research Center in Denver, Colorado, also recognized the need to identify specific behavioral processes to improve both accuracy and training time and then elicit these processes into clear and concise instructions that are compatible with the skills and abilities of the users.

Task analysis was widely adopted and matured by a variety of practitioners, but behavioral cues and the role of behavioral processes were not always used due to inconsistent applications of the technique as new task analysis methodologies were developed (Davies, 1973; Stammers et al., 1990). For example, Mager (1962) renamed the stimulus elements as "conditions" and this resulted in the stimulus term often being misinterpreted as environmental (e.g. temperature) or setup conditions to the experiment; as a result the *behavioral* stimulus elements were often lost (Shriver et al., 1982).

In 1969 the Special Interest Group on Task Analysis of the Human Factors Society (chaired by Shriver) counted over 80 variations of task analysis in use at the time; in many instances the behavioral cues were missing. This led to the sub-categorization of task analysis methods into more specific methodologies such as basic task analysis (Miller, 1953), job analysis (Chapanis, 1970), functional analysis (Shriver, 1960), and link analysis (Thomson, 1972). A major output of the interest group's effort was the establishment of the behavioral task analysis (BTA) technique that returned the stimulus element to task analysis. The BTA technique was later codified in a U.S. Air Force technical report (Shriver, 1975) that would also serve as an important milestone in the development of JPAs.

The first research activities that specifically used the term "job performance aid" began with more U.S. military contracts in the late 1950s and early 1960s in conjunction with task analysis

(e.g. Shriver, Fink, & Trexler, 1959; Folley, 1961; Folley & Shettel, 1962; Goff et al., 1969). Over the next two decades JPAs were being studied in a variety of contexts. The early military studies tended to look for ways to reduce errors in complex tasks that were not often performed correctly (such as repair and maintenance activities) or to shorten the training time necessary for novice technicians to acquire such skills (Gebhard, 1970; Elliott & Joyce, 1971; Foley, 1972). Later studies were refined to determine how different JPA formats (pictures or printed text) could convey information differently and were therefore better suited for particular circumstances (Booher, 1975; Smillie, 1978; Smillie & Ayoub, 1980). Throughout these experiments the JPA was found to be an effective tool at simplifying tasks for novice users that would normally require extensive training or complex information processing, and errors in human performance were reduced significantly. These experiences led to a variety of specifications (Folley et al., 1971; Joyce et al., 1973a; Shriver, 1975) and handbooks (Joyce et al., 1973b; Booher, 1978; Smillie, 1985) to assist developers in creating effective JPAs for a wide variety of tasks in both military and commercial industries.

Meanwhile, specific industries were conducting their own studies on the limitations of human performance. McKenzie (1958) was an early proponent of having clear standards and instructions in order to improve the accuracy and consistency of inspectors. This followed with more than three decades of research from the industrial inspection community (summarized in detail by Fox, 1973; Wiener, 1975; Drury, 1982, and See, 2012), with some of it focused on vigilance theory (e.g. Elliott, 1960; McGrath and Harabedian, 1961; Mackie, 1964; Baddeley and Colquhoun, 1969; Tsao, Drury, and Morawski, 1979; Tsao and Wang, 1984; Murray and Caldwell, 1996) that stemmed from the pioneering work of Mackworth (1950). These studies were paralleled by researchers who began to understand the underlying models of human information processing that would become useful to JPA developers (Gagne, 1962; Kibler, 1965; Harris, 1969; Norman, 1981; Rasmussen, 1982; Norman, 1983; Rasmussen and Vicente, 1989; Reason, 1990; Barshi and Healy, 1993; Allen and Rankin, 1997).

Following the Three Mile Island accident (Kemeny, 1979) the nuclear power industry commissioned a series of studies to understand and eliminate the potential for human errors, and JPA research experienced a resurgence of interest (Clark, 1982; Shriver et al., 1982; Mattson, 1989; Hallbert et al., 1992). In the years since, other high consequence industries such as commercial aviation and medicine have followed suit by embracing a JPA focused approach (most notably checklists) as well as wider principles of error management, human factors, and quality assurance (Helmreich and Schaefer, 1994; Helmreich, 2000; Helmreich and Merritt, 2001; Haynes et al., 2009; Gawande, 2010; White et al., 2010).

SNL designs and builds a variety of high consequence products, however its primary mission is with nuclear weapons. SNL is therefore subject to the guidance and oversight by the Department of Energy (DOE). As reviewed in Stevens-Adams (2013), the DOE *Human Performance Handbook* (DOE, 2009) focuses on tools and methods for preventing and mitigating human error, including processes specifically related to verification activities that are useful for quality assurance:

- **Self-Checking** – A process by which the performer pauses to reflect on the intended action and its expected outcome, whether the proposed action is correct for the situation,

and resolve any questions or concerns *before* proceeding by using the STAR (Stop, Think, Act, Review) practice.

- **Peer Checking** – A series of actions by 2 individuals working together at the same time and place, before and during a specific action, as an *augmentation but not a replacement* for self-checking.
- **Concurrent Verification** – A series of actions by 2 individuals, at the same time and place, to separately confirm the condition before, during, and after an action, when the consequences of an incorrect condition would subsequently lead to undesired harm.
- **Independent Verification** – A process by which one individual, separated by time and distance from the action, confirms the desired condition.
- **Peer Review** – A process by which qualified reviewers read and check the quality of another’s work, and is an informal technique that does not supplant required procedures or checklists.

As further discussed in Stevens-Adams (2013), DOE Order 422.1, *Conduct of Operations*, requires independent and concurrent dual verification techniques to be implemented for critical equipment configurations. These two techniques are further defined in the *DOE Guide to Good Practices for Independent Verification* (DOE, 1993):

- **Independent Verification**–The act of checking, by a separate qualified person, that a given operation or the position of a component conforms to established criteria.
- **Concurrent Dual Verification**–A method of checking an operation, an act of positioning, or a calculation in which the verifier independently observes and/or confirms the activity.

There has been a divergence between the term *inspector*, which implies an independent reviewer of a particular item or task, and *checker* that is specified by the verification techniques. As a result of the above definitions (plus the fact that this study is about *check-lists*), the term *checker* will be used throughout this paper to refer to the QA role being performed by the experimental participants.

The motivation for this study should now be understood as an intersection between the complementary disciplines of human factors and QA (itself derived from the tradition of inspection), with a focus on the concurrent dual verification activity within high consequence environments. The JPA approach continues a contemporary interest in checklists while leveraging the legacy techniques of task analysis, plus it is one of many tools that a QA practitioner can use as a control against the risk of defects. The concurrent dual verification JPA thus fills a research void and should draw interest from a variety of disciplines.

This Page Intentionally Left Blank

### 3. METHODOLOGY

This chapter describes the methodology and approach used for the research study. Section 3.1 specifies the approach behind the selection of the experimental task and presents an appropriate and novel test apparatus as a contribution to the wider body of knowledge. Section 3.2 explains the experimental design. Section 3.3 explains how the number of experimental participants was determined and Section 3.4 discusses the selection of participants and the presentation of the experiment. Section 3.5 details the rationale and configuration of the job performance aid used in this study.

#### 3.1 Experimental Task Selection

The primary emphasis of the original research question is on the effectiveness of a JPA on a concurrent dual verification task. However, there must be a task for the QA checkers to observe such that they can perform concurrent dual verification. Basic guidelines for selecting this task were as follows:

- Not be too complex or time-consuming
- Not be too simple, such that the ability to inject faults would be difficult as they would be too obvious
- Be consistent with a task that may be conducted in a high consequence manufacturing environment

A predefined Lego™ assembly task (with visual assembly instructions) was determined to satisfy these conditions, with a corresponding checklist for the QA checker to follow (or have absent in the control condition). An advantage of selecting a Lego™ assembly task over more applied techniques (such as repairing a lawnmower engine or replacing a circuit board in a personal computer) is that there is not any built-in covariance of prior knowledge for experimental participants who may have expertise in these applied tasks. In other words, selecting simple tasks would reduce uncontrolled differences in participant expertise. Another advantage is the similarity of the Lego™ task with those conducted in a high consequence manufacturing environment. Assemblers are often provided with a kit of similar-looking parts (e.g., fasteners of different size but equal length) and instructions how to install them, and in many real-world applications there is a QA checker to oversee the task and ensure that no errors are made.

A within-subjects design model was chosen in order for participants to have their performance evaluated both with and without the checklist, so this drove the need for a second assembly task. Due to the potential for learning effects, two *different* Lego™ assembly tasks (Pattern A, Pattern B) were identified in which faults could be introduced by the inclusion of incorrect pieces into the patterns. Another variable to be reasonably controlled is the difference in size or complexity between the two patterns. This constraint eliminated many of the popular Lego™ patterns from contention, such as Lego™ City or themed sets like Star Wars™ or super heroes, since the ability to equalize both size and complexity across two patterns is difficult. Further, defined patterns with uniquely shaped pieces typically require all parts to be used during assembly (i.e., there are no pieces left over). This might make it too obvious to inject faults into these complex

patterns. The option chosen for this experimental task was to select two similar but existing patterns (so the instructions were already developed and available) and then combine them into a larger boxed kit (thus each pattern was a sub-assembly within the larger kit).

Pattern A has 104 pieces with 7 of them specifically chosen for inserting a fault. Pattern B has 150 pieces, with another 7 pieces specifically chosen for inserting a fault. Since one of the constraints is to avoid selecting a task that is too simple, the patterns were not assembled as stand-alone kits but instead contained within a larger set of pieces that may or may not be used in either sub-assembly. This eliminates the potential confound where the QA checker might notice that an “incorrect” part was used if there are supposed to be no pieces remaining when the assembly task is completed. The total number of available pieces needs to be higher than both sub-assemblies combined, with additional margin such that there is sufficient uncertainty in the experiment to warrant the use of a JPA. Since the total number of pieces in Pattern A and Pattern B is 254 (almost  $256 = 2^7$  or 7 bits of information), a total assembly kit of 512 pieces was chosen since  $512 = 2^9$  or 9 bits of potential information (Posner, 1964).

However, the total number of bits of information is much higher than 9 bits when considering the different part shapes (23 for Pattern A), colors (6 for Pattern A), and markings that each act as multipliers to the total number of possible combinations of parts within the two patterns. This large number is roughly doubled when adding in Pattern B, and doubled again when considering parts not used in either pattern. Thus there is tremendous (and sufficient) uncertainty in the kit of 512 parts that experimental participants were not able to determine the total number of pieces being used for each pattern. If the JPA effectively improves the chances that the *correct piece* in the *correct color* and *correct shape* is installed in the *correct way*, and it is “infused” (Oberhettinger, 2012) at the *correct time* and *correct place* in the assembly process, then it would be a tremendously useful tool to support concurrent dual verification in mitigating human error.

In order to more fully describe the correct piece / correct place discussion above, consider that each of the 512 pieces belongs to only one of the following five part groups:

- Parts used only in Pattern A: 52 pieces
- Parts used only in Pattern B: 98 pieces
- Parts that exist in both patterns A and B: 104 pieces (or 52 in both patterns)
- “Incorrect” parts that are used to interject faults into both patterns: 9 pieces
- Parts not used in either pattern: 249 pieces

To further clarify, Pattern A has 52 pieces that are unique and 52 that are also used in Pattern B, for a total of 104 pieces. Pattern B has 98 pieces that are unique and 52 that are also used in Pattern A, for a total of 150 pieces. Mutual exclusivity of the five part groups allows exact probabilities to be determined for analysis and compared to experimental results. The inclusion of both patterns in the kit of 512 pieces, as well as the decision to not return the first pattern’s pieces to the kit before commencing with the second pattern in the trial, are counterbalanced by the randomized assignment of participants to one of the four test conditions.

### 3.2 Test Approach

In the Lego™ assembly task described in Section 3.1, experimental participants performed the role of QA checker and the author acted as the assembler of the patterns. A distinction was made between an **error** by the assembler that was purposely made (**fault**) in order to elicit experimental data and an error that was unintentionally made (**mistake**) and is less germane to the results. However, the use of the term “error” was necessary in providing instructions to the checkers because they correctly identified both faults and mistakes during the experiment without any knowledge of the difference between the two. The instructions given to the participant before the test began were: “Please let me know if you notice any errors. “ When a fault was made and the QA checker did not notice, this was defined as a **miss**; if they did notice, it was defined as a **detection**. Some faults were detected immediately, but others were detected later in the assembly process. This difference was anticipated and considered important to the results, so there needed to be a way to distinguish between the two responses. While the identification of a fault before turning the page in the assembly instructions is a detection, if the fault was identified after the page was turned it was determined to be a **catch**. A **false alarm** occurred when an error (fault or mistake) was not performed but the QA checker notified the assembler that they thought one had occurred. If the assembler intentionally disguised a fault from the QA checker then it was classified as a **violation** (though the experiment was designed to eliminate the risk of violations and none were present in the results). Finally, **sequence errors** are a special type of false alarm that is similar to Norman’s (1981) mode error, where the response is incorrect and represents a user assuming a state other than the true state of the task. These definitions, specific to this study and not necessarily in common use elsewhere, are summarized in Table 1.

**Table 1: Definitions Used in the Assembly Task**

Fault	An intentional and specific error that is introduced by the assembler
Mistake	An unintentional error by the assembler
Error	The use of an incorrect piece or incorrect assembly order
Miss	An error that is not noticed by the QA checker; includes catches
Detection	Before turning the next page of the assembly instructions, the QA checker identifies an error
Catch	An error that is noticed by the QA checker later (after the page has been turned)
False alarm	Any response from the QA checker when there is no error present
Violation	A fault that is purposely hidden from the QA checker, such that there is not an opportunity to notice it
Sequence error	The QA checker verbally authorizes the assembler to turn the page of the instructions before the assembly task on that page has been completed; treated as a false alarm

Faults were inserted by the assembler into both patterns in specific pre-determined locations, with the pre-test instruction that the QA checker observe the assembly task and tell the assembler if they noticed any errors. The fault types are:

1. *Marking Fault*: Insert the incorrect piece that has markings. For example, instead of inserting the “correct” 2 x 2 yellow piece another 2 x 2 yellow piece with a pattern on one side is installed.
2. *Incorrect Piece*: Insert the incorrect piece(s) but with no markings. For example, instead of inserting the “correct” 2 x 4 black piece two 1 x 4 black pieces is installed.
3. *Wrong Order*: Insert the correct piece(s) but in the incorrect configuration. For example, instead of constructing a 2 x 10 wall with a 2 x 8 piece on the left and a 2 x 2 piece on the right, the order is switched and the 2 x 2 piece goes on the left
4. *Wrong Orientation*: Insert the correct piece(s) but in the wrong location or orientation. For example, a window piece is installed backwards, or offset by one row from what the “correct” orientation is.

The fault types were chosen for their reasonable similarity to those that occur in high consequence environments. For example, fault type 4 (wrong orientation) is similar to a component being installed backwards on a printed circuit board. The author has past experience with fault type 1, dating from an incident whereby a damaged part with noticeable indications of wear was nearly installed onto a larger assembly. The specific faults in the Lego™ assembly task were only selected for their ease of insertion and ability to avoid detection by the QA checker, and were not equally spaced throughout the assembly task. These faults are further detailed in Section 3.4.

The independent variables are as follows:

- Pattern order (A first or B first)
- JPA (present or not present)

Uncontrolled variables (potential covariants) that were observed are:

- Pattern A elapsed time
- Pattern B elapsed time
- Patterns A and B combined elapsed time
- Fault type

Because the errors (both faults and mistakes) are performed by the assembler they are independent of the presence or absence of the JPA. Therefore the responses from the participant (the QA checker) represent the primary dependent variables for this study. For purposes of simplifying the statistical design the dependent variables have been categorized as described above and summarized into only 3 response levels: 1) detections; 2) misses, and 3) false alarms.

The primary purpose of this experiment is to determine whether or not the JPA has an effect on the performance of the QA checker. Therefore, the first null hypothesis to be tested is:

H<sub>1</sub>: The presence of a JPA has no effect on the detection of faults in the quality assurance role.

with the alternative hypothesis being:

H<sub>1A</sub>: The presence of a JPA **has** an effect on the detection of errors in the quality assurance role.

The following additional hypotheses are being tested in this experiment, specifically addressing the independent variables listed previously:

H<sub>2</sub>: There is no difference in the detection of faults with a JPA than without a JPA.

H<sub>2A</sub>: There **is** a difference in the detection of faults with a JPA than without a JPA.

H<sub>3</sub>: The order of presentation of the JPA has no effect on the detection of faults in the quality assurance role.

H<sub>3A</sub>: The order of presentation of the JPA **has** an effect on the detection of faults in the quality assurance role.

H<sub>4</sub>: The order of presentation of the different patterns has no effect on the detection of faults in the quality assurance role.

H<sub>4A</sub>: The order of presentation of the different patterns **has** an effect on the detection of faults in the quality assurance role.

An additional hypothesis, H<sub>5</sub>, is explained in the next section.

### 3.3 Sample Size Determination

The minimum number of participants was determined from the estimated probability of detecting each fault, with different values estimated for with and without a JPA. The different probabilities of detection among the 14 total faults span a fuzzy range between somewhat-easy and somewhat-difficult, and the added complexity of 4 different fault types with uneven placement throughout the two patterns made it necessary to create simplifying assumptions for both easier statistical analyses and a testable hypothesis. The probabilistic Rasch (1960) model, which in turn requires an ordered index based on the Guttman (1944) scale, is an appropriate guide for this experiment as it models the trade-off between the QA checker's abilities and the difficulty of detecting each fault.

One simplifying assumption was to assign a constant average probability of detection across all fault instances. Since the checklist specifically identified the 4 different fault types, it was reasonable to assume that the average probability of detection is  $p_1=0.5$  without a JPA and  $p_2=0.9$  with a JPA. The number of experimental participants (n) must be a multiple of 4 to match

the number of test conditions (listed in Section 3.1) and therefore make the experiment balanced. For  $n*14$  binary trials ( $n*7$  with a JPA and  $n*7$  without a JPA), the probability of concluding that there is a difference when  $p_1=p_2$  depends on  $n$  as follows in Table 2:

**Table 2: Probability of Correctly Concluding that  $p_2 > p_1$**

<b>n</b>	<b>Type-1 error of 0.05</b>	<b>Type-1 error of 0.025</b>
4	0.96	0.93
8	0.9995	0.998
12	~1	0.99998
$\geq 16$		~1

This analysis suggests a minimum sample size of eight participants. However, the simplifying assumption of average probability of detection across the individual 14 faults suggests that this number should be higher. Once the experimental trials began it soon became apparent that the addition of more participants was not particularly onerous on the assembler or support staff and could be accomplished in a reasonably short timeframe. Thus the number of participants for this experiment was raised to 24. In addition, the assumption of an average probability of detection being uniformly distributed across the experiment is an additional hypothesis to be tested:

$H_5$ : The probability of detection for each fault is equal.

$H_{5A}$ : The probability of detection for each fault is **not** equal.

### 3.4 Test Procedure

Participants consisted of permanent and contract employees (14 males, 10 females) from SNL. Candidates who perform QA activities as part of their normal job responsibilities were specifically excluded from this study, so all QA checkers were considered novices in this role. The experiment was conducted in a video recording studio to facilitate both the capture of the assembly tasks and audio of the QA checker's interactions with the assembler. The video camera was placed such that all responses were recorded but not the faces of participants, in order to maintain confidentiality of participation. Figure 1 shows the test setup from one of the experimental trials with the assembler on the left and the QA checker seated off-camera on the right.



**Figure 1. Test Setup**

The assembly instructions were printed and used by the assembler in clear view of the QA checker, along with the collection of 512 pieces. These assembly instructions were the original instructions for each pattern developed by Lego™, with no changes or additions. The pre-test instructions given verbally to the QA checkers (only once, before the first pattern was presented) were to observe the assembly task and report any errors as soon as they were noticed, and not to touch any of the Lego™ pieces or assist the assembler in locating them. Participant #7 was disqualified due to misunderstanding the pre-test instructions. Another participant was recruited to replace #7 prior to the completion of the study, maintaining a total of 24 valid participants. For one of the two assembly tasks in each experimental trial, the checklist was provided for review beforehand and was available throughout the duration of that assembly task. Table 3 summarizes how the order of patterns varied across the 24 participants. By varying the order of assembly in this way, specific effects of the presence of a JPA (if they exist) can be distinguished from the effects of sequence of assembly ( $H_2$ – $H_3$ ).

**Table 3: Grouping of Experimental Participants By Sequence of Assembly**

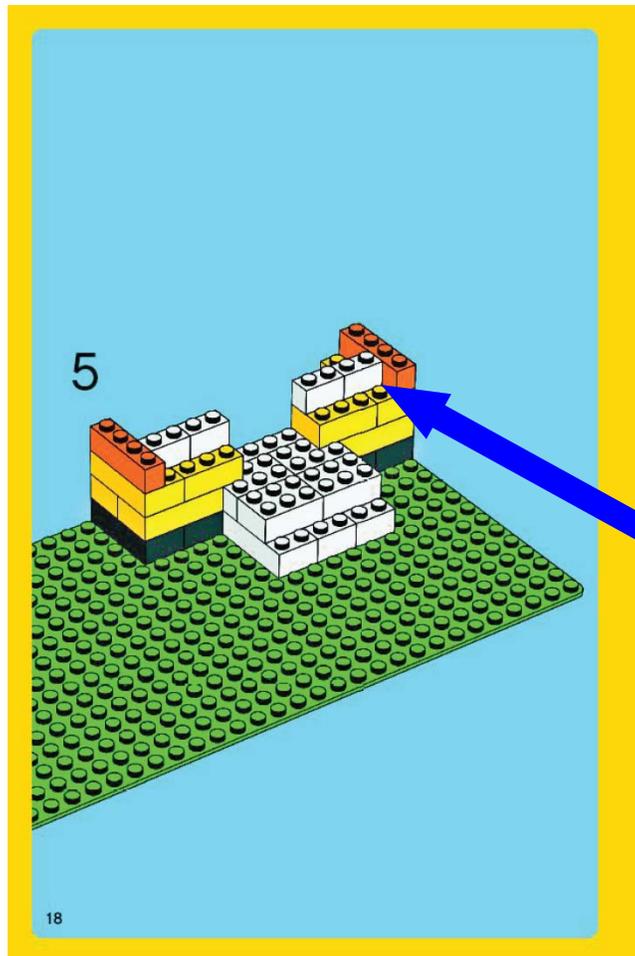
Number of Participants	Sequence of Assembly, Presence of JPA	Abbreviation
6	Pattern A without JPA, followed by Pattern B with JPA	A{JB}
6	Pattern A with JPA, followed by Pattern B without JPA	{JA}B
6	Pattern B without JPA, followed by Pattern A with JPA	B{JA}
6	Pattern B with JPA, followed by Pattern A without JPA	{JB}A

As explained earlier in Section 3.2, the assembler intentionally inserted faults in 14 specific locations (seven in each pattern) and there were four different types of faults. Table 4 summarizes the specific faults and associated fault types.

**Table 4. Fault Descriptions/Types**

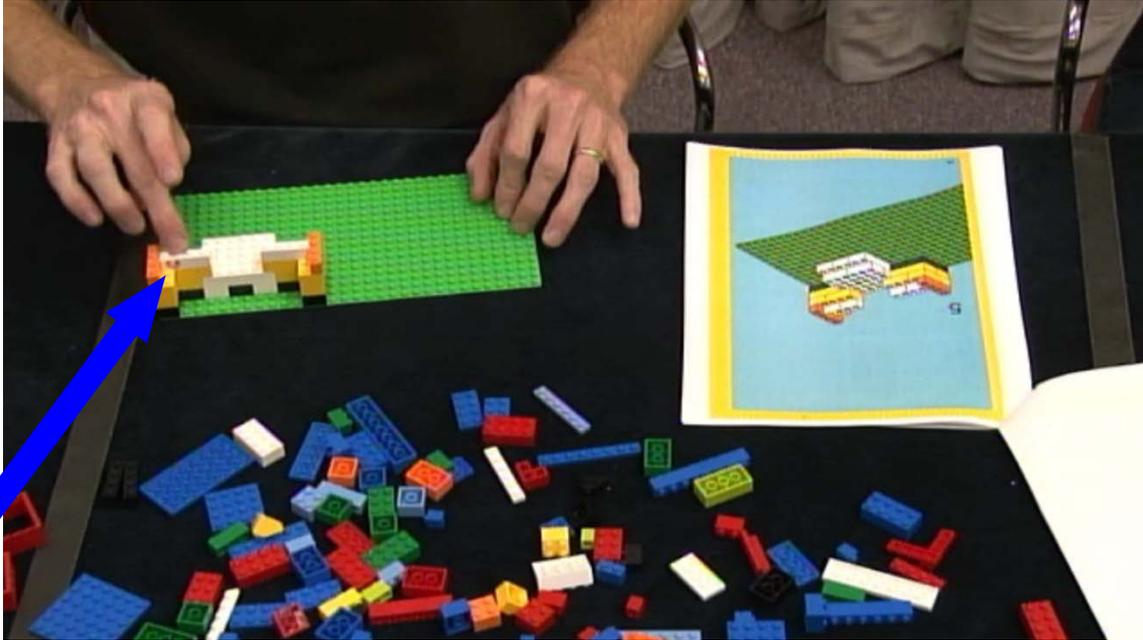
<b>Fault</b>	<b>Fault Description</b>	<b>Fault Type</b>
Fault #1:	Inject kitty 1x2 white piece on far right, facing backward (p.5)	1
Fault #2:	Switch order of 1x1 and 1x3 yellow pieces on left side (p.8)	3
Fault #3:	Switch order of 2x2 and 2x4 yellow pieces on right side (p.9)	3
Fault #4:	Switch order of 1x1 and 1x2 orange pieces on right side (p.14)	3
Fault #5:	Install window backwards (p.15)	4
Fault #6:	Inject kitty 1x2 white piece, facing forward (P.19)	1
Fault #7:	Inject 1x1 eye piece, facing right (P.20)	1
Fault #8:	Make 4th 2x4 black block two 1x4s (p.2)	2
Fault #9:	Don't install red window ahead of white door (p.6)	4
Fault #10:	Inject 1x2 yellow kitty on right side, facing forward (p.10)	1
Fault #11:	Switch order of 2x2 and 2x4 white pieces on right side (p.12)	3
Fault #12:	Inject kitty 1x2 white piece on right side, facing left -- towards QA (p.14)	1
Fault #13:	Make two 1x4 white blocks one 2x4 block (p.14)	2
Fault #14:	Inject 2x2 yellow eye block, face forward (p.16)	1
		1 = Markings
		2 = Incorrect piece(s)
		3 = Wrong order
		4 = Wrong orientation

Figure 2 shows the specific page from the assembly instructions during which Fault #1 was inserted; fault #1 was also fault type 1. The arrow points to the white 1 X 2 Lego™ piece that was chosen for replacement with one of identical shape and color, except with markings on one side.



**Figure 2. Sample Instructions Page, Depicting Fault #1 (Fault Type 1)**

The actual replacement of this piece in an experimental trial is shown by the arrow in Figure 3, where markings are shown on the back side of the part that is facing away from the QA checker. During experimental design it was thought that pieces with markings would be too easily noticeable if they were facing the checker, so it was determined that most of them would face in a different direction. As a counterbalance against these faults then becoming too *difficult* to notice, all pieces with markings were placed face-up on the table during the trial such that the QA checker would have the *opportunity* to notice them.



**Figure 3. Fault #1 Occurring in an Experimental Trial**

All four fault types are illustrated in Appendix B.

### **3.5 Job Performance Aid**

This experiment was designed to focus not on the Lego™ assembly task itself but rather the effectiveness of a JPA for the experimental participant that is performing the role of the QA checker. Prior JPA research (Miller, 1956; Gebhard, 1970; Elliott & Joyce, 1971; Foley, 1972; Clark, 1982; Shriver et al., 1982, Barshi & Healy, 1993) has studied the use of JPAs with both novice and expert users, but there is general agreement that they are most effective as training aids for novice users. Some common themes in the literature proved useful in the development of the checklist for this study:

- **The focus is on the user.** JPA developers should use task analysis techniques to fully understand the job function and the behaviors used when the task is performed (Gebhard, 1970; Clark, 1982; Shriver et al., 1982; Smillie, 1985).
- **Information must be task oriented.** The JPA should identify exactly what the user needs to do using brief, concise, and explicit instructions to aid short term memory. Instructions should be directive and action-specific using simplified and standard language, and contain only information that is necessary (McKenzie, 1958; Gebhard, 1970; Clark, 1982; Shriver et al., 1982; Smillie, 1985; Hallbert et al., 1992; Gawande, 2010)

The JPA for this experiment (see Appendix A) consists of a short, concise, and simple checklist intended to elicit behaviors that would enhance the detection of faults in this experiment. Each of the six checklist items was considered to be essential information, with the following rationale:

- *Your role as an observer is an essential part of this important task. Complex assemblies require a second set of eyes in order to catch any errors.* It is the author's experience that novice QA checkers do not always recognize the importance of simply being an observer for an important task. McKenzie (1958) noted that inspection takes place in a social context where there may be different 'social norms' of rejection, which is consistent with the author noting that subject matter experts who perform the work may sometimes resent a non-expert "checking their work." In addition, Wiener (1975) reminds us of Herzberg's (1966) theory of industrial motivation and quotes Lucaccini et al. (1968) where subjects receiving more positively-toned instructions detected more signals.
- *Pay attention for the following types of error.* Both Clark (1982) and White et al. (2010) recommend that hazards or high-risk concerns have their own specific checklist items, though to avoid confusion the four fault types were addressed in only three statements.
- *Feel free to ask questions about the task at any time. If necessary, ask the assembler to stop until you are comfortable with proceeding.* This statement was intended to build the QA checker's confidence by signaling that they are vital to the assembly process and could stop it at any time. It also grants permission to talk to the assembler and not be considered a distraction.
- *The assembler should not turn to the next page of the instructions without your approval.* This statement was necessary to establish the event horizon at which a detection becomes a catch. The specific emphasis on catches was kept from the QA checker as it would have introduced a confound.
- *For each page of the instructions, the order of assembly does not matter.* To avoid confusion that could have led to an increase in false alarms, this statement needed to be explicit.
- *The box contains 512 total parts. Some parts will be used and some will not.* This statement eliminates the confound whereby the QA checker could use the presence or absence of parts on the table as a cue.

Clark (1982), Shriver et al. (1982), and Smillie (1985) are emphatic that the key final step in successful development of a JPA is verification and validation with expert users. The process relies heavily on these reviews to identify and correct procedural ambiguities, omissions, and inaccuracies, and depending on the complexity of the JPA the process may require multiple iterations. For the simple and concise checklist in this experiment, the author conducted a pilot study with four QA experts in the Surety Engineering Group at SNL for additional guidance and a final verification of the checklist format. The study was approved (before the pilot study) by both the SNL and Purdue University Institutional Review Boards under protocols SNL1349 and IRB00001150, respectively.

This Page Intentionally Left Blank

## 4. RESULTS AND ANALYSIS

Binary responses for 14 fault observations were recorded for each participant, and are summarized in Table 5. Participant #7 was disqualified due to misunderstanding the pre-test instructions and was replaced with participant #25. The three Pattern B instances where only 6 trials appeared were from incorrect assembly and the planned fault could not be inserted. Note that participant #1 detected every fault in both patterns.

**Table 5: Fault Detection Performance, By Participant**

Subject #	Pattern A Trials	Pattern A Detections	Pattern B Trials	Pattern B Detections	Percent Detected
1	7	7	7	7	100%
2	7	5	6	5	77%
3	7	4	7	3	50%
4	7	7	6	3	77%
5	7	6	6	4	77%
6	7	5	7	4	64%
8	7	4	7	5	64%
9	7	6	7	5	79%
10	7	5	7	7	86%
11	7	3	7	4	50%
12	7	3	7	4	50%
13	7	3	7	4	50%
14	7	3	7	6	64%
15	7	4	7	4	57%
16	7	4	7	5	64%
17	7	6	7	3	64%
18	7	3	7	4	50%
19	7	4	7	2	43%
20	7	4	7	3	50%
21	7	4	7	5	64%
22	7	4	7	5	64%
23	7	4	7	5	64%
24	7	3	7	3	43%
25	7	4	7	3	50%

Table 6 summarizes the results for each particular fault. The three instances of incorrect assembly are seen again where two of the trials associated with fault #8 and one trial with fault #11 were deemed to be a “no test.” It is noteworthy that fault #2, fault #4, and fault #11 were always detected while fault #1, fault #6, fault #7, fault #10, and fault #14 were frequently missed. The members of the latter group (#1, #6, #7, #10, and #14) are all marking faults (fault type 1); however, fault #12 (also fault type 1) was specifically designed to be noticed and, as expected, frequently detected<sup>2</sup>.

---

<sup>2</sup> This intentionally-noticeable fault was designed to address the vigilance decrement noted by Tsao, Drury, and Morawski (1984) where a subject’s expectations of finding another defect appeared to decrease immediately after a defect had been found. However, hypotheses related to vigilance theory were later dropped from this study.

**Table 6: Performance, By Fault Number**

Pattern	Fault Number	Fault Type	Number of Trials	Number of Detects	Percent Detected
A	1	1	24	5	21%
A	2	3	24	24	100%
A	3	3	24	23	96%
A	4	3	24	24	100%
A	5	4	24	17	71%
A	6	1	24	6	25%
A	7	1	24	6	25%
B	8	2	22	15	68%
B	9	4	24	21	88%
B	10	1	24	5	21%
B	11	3	23	23	100%
B	12	1	24	20	83%
B	13	2	24	17	71%
B	14	1	24	2	8%

Using the performance summaries in Table 5 and Table 6 as a guide, binary logistic regression (Agresti, 2013) was used to model the probability of detecting a fault as a function of the experimental factors. In logistic regression, the probability of detection,  $\pi$ , is assumed to relate to the experimental factors  $(x_1, x_2, \dots, x_p)$  via the general expression  $\log \left( \frac{\pi(x_1, x_2, \dots, x_p)}{1 - \pi(x_1, x_2, \dots, x_p)} \right) = f(x_1, x_2, \dots, x_p)$ , where  $p$  is the number of factors. Note that  $\log \left( \frac{\pi(x_1, x_2, \dots, x_p)}{1 - \pi(x_1, x_2, \dots, x_p)} \right)$  is commonly referred to as the *logit* or log odds. Here, we further assume that  $f$  is a linear combination of the factors. Collectively, the set of values for the experimental factors describes the experimental condition.

Due to the nature of the experiment and the apparent differences in detection ability across faults, separate probability models were developed for each pattern. For Pattern A, faults #2 and #4 were excluded from the models because they were always detected and therefore there is no ability to discern an effect of sequence in those cases; for Pattern B, fault #11 was excluded for the same reason. The model form used in each case is:

$$\log \left( \frac{\pi(\text{Err}(i), \text{Seq}(j))}{1 - \pi(\text{Err}(i), \text{Seq}(j))} \right) = \alpha_0 + \beta_i + \gamma_j, \quad (1)$$

where  $\alpha_0$  represents the log odds at a standard experimental condition (defined by a standard error number and a standard sequence),  $\beta_i$  reflects the change in log odds when changing the experimental condition from the standard fault number to fault # $i$ , and  $\gamma_j$  reflects the change in log odds when changing the experimental condition from the standard sequence to sequence  $j$ . In the case of both patterns, the standard sequence is denoted by pattern A, followed by the JPA instructions and then pattern B, or in study notation, A{JB} (see Table 3). In the case of pattern A, fault #1 is the standard fault number and in the case of pattern B, fault #8 is the standard fault

number. The form of the model allows for an assessment of the effect of the JPA given differences in detection ability across fault numbers. The effect of fault number is compensated for by the  $\beta_i$  terms, while the effect of the JPA is associated with the  $\gamma_j$  terms (more on this in Chapter 5).

Minitab (version 16.1.1) was used to estimate the model parameters for each pattern. Table 7 and Table 8 display the parameter estimates (and associated *p-values*) related with patterns A and B, respectively. The *p-values* measure whether or not the model terms are significantly ( $p < 0.05$ ) different than zero.

**Table 7: Logistic Regression Table, Pattern A**

Parameter	Estimate	SE Estimate	Z-ratio	P-value
$\alpha_0$	-2.845	0.810	-3.51	0.000
$\gamma_{B\{JA\}}$	1.792	0.776	2.31	0.021
$\gamma_{\{JA\}B}$	1.999	0.778	2.57	0.010
$\gamma_{\{JB\}A}$	1.578	0.775	2.04	0.042
$\beta_3$	4.967	1.218	4.08	0.000
$\beta_5$	2.494	0.731	3.41	0.001
$\beta_6$	0.251	0.710	0.35	0.724
$\beta_7$	0.251	0.710	0.35	0.724

**Table 8: Logistic Regression Table, Pattern B**

Parameter	Estimate	SE Estimate	Z-ratio	P-value
$\alpha_0$	0.893	0.593	1.50	0.132
$\gamma_{B\{JA\}}$	0.350	0.637	0.55	0.582
$\gamma_{\{JA\}B}$	-0.615	0.610	-1.01	0.313
$\gamma_{\{JB\}A}$	-0.187	0.612	-0.31	0.760
$\beta_9$	1.211	0.776	1.56	0.119
$\beta_{10}$	-2.150	0.691	-3.11	0.002
$\beta_{12}$	0.870	0.722	1.20	0.228
$\beta_{13}$	0.132	0.650	0.20	0.839
$\beta_{14}$	-3.227	0.879	-3.67	0.000

For the various experimental conditions, the fitted model can be used to estimate the probability of detection by inverting equation (1) and replacing the parameters with their estimates. For example, suppose that we are interested in the probability of detecting fault #1 (pattern A) assuming the standard sequence A{JB}. Then,  $\hat{\pi}(Err(1), Seq(A\{JB\})) = \frac{\exp(\hat{\alpha}_0)}{1+\exp(\hat{\alpha}_0)} = 0.0549$ , where  $\hat{\alpha}_0 = -2.845$  (from Table 7). Further suppose that we are interested in the probability of detecting fault #6 (pattern A) assuming the sequence B{JA}. Then,  $\hat{\pi}(Err(6), Seq(B\{JA\})) = \frac{\exp(\hat{\alpha}_0 + \hat{\beta}_6 + \hat{\gamma}_{B\{JA\}})}{1+\exp(\hat{\alpha}_0 + \hat{\beta}_6 + \hat{\gamma}_{B\{JA\})} = 0.3095$ , where  $\hat{\alpha}_0$ ,  $\hat{\beta}_6$ , and  $\hat{\gamma}_{B\{JA\}}$  are from Table 7. Table 9 and Table 10 list the estimated probabilities of detection (for all experimental conditions) and compare with the observed fraction of faults detected for patterns A and B, respectively. Note that the estimated probability of detection matches reasonably well with the observed fraction detected

for all cases. The implication is that the predictors contained in the equation (1) model contain all of the necessary information, meaning that this is an accurate model for the experimental conditions. The same conclusion was reached from Minitab using the Pearson ( $p=0.171$ ), Deviance ( $p=0.194$ ), and especially Hosmer-Lemeshow ( $p=0.725$ ) formal tests where p-values *greater* than 0.05 indicate that there is no evidence for lack-of-fit and that the model is accurate. These results would not be statistically different from increasing the pool of participants except in the case of a *much* larger number that would be required to detect small differences.

**Table 9: Estimated Probability of Detection, Pattern A**

Fault #	Sequence	Estimated Probability of Detection	Observed Fraction Detected
1	A {JB}	0.055	0.000
1	B {JA}	0.259	0.500
1	{JA} B	0.300	0.167
1	{JB} A	0.220	0.167
3	A {JB}	0.893	1.000
3	B {JA}	0.980	1.000
3	{JA} B	0.984	1.000
3	{JB} A	0.976	0.833
5	A {JB}	0.413	0.500
5	B {JA}	0.809	0.500
5	{JA} B	0.839	0.833
5	{JB} A	0.773	1.000
6	A {JB}	0.069	0.000
6	B {JA}	0.310	0.333
6	{JA} B	0.355	0.500
6	{JB} A	0.266	0.167
7	A {JB}	0.069	0.000
7	B {JA}	0.310	0.333
7	{JA} B	0.355	0.333
7	{JB} A	0.266	0.333

**Table 10: Estimated Probability of Detection, Pattern B**

<b>Fault #</b>	<b>Sequence</b>	<b>Estimated Probability of Detection</b>	<b>Observed Fraction Detected</b>
8	A {JB}	0.709	0.833
8	B {JA}	0.776	0.800
8	{JA} B	0.569	0.400
8	{JB} A	0.670	0.667
9	A {JB}	0.891	0.833
9	B {JA}	0.921	0.833
9	{JA} B	0.816	0.833
9	{JB} A	0.872	1.000
10	A {JB}	0.221	0.167
10	B {JA}	0.288	0.333
10	{JA} B	0.133	0.333
10	{JB} A	0.191	0.000
12	A {JB}	0.854	1.000
12	B {JA}	0.892	0.667
12	{JA} B	0.759	0.833
12	{JB} A	0.829	0.833
13	A {JB}	0.736	0.667
13	B {JA}	0.798	0.833
13	{JA} B	0.601	0.500
13	{JB} A	0.698	0.833
14	A {JB}	0.088	0.000
14	B {JA}	0.121	0.333
14	{JA} B	0.050	0.000

This Page Intentionally Left Blank

## 5. DISCUSSION

The effect of the JPA can be deduced by considering the complete set of estimated  $\gamma_j$  terms in the model, which represent the 4 test conditions shown in Table 3. In the case of pattern A (see Table 7), note that  $\hat{\gamma}_{B\{JA\}}$ ,  $\hat{\gamma}_{\{JA\}B}$ , and  $\hat{\gamma}_{\{JB\}A}$  are all statistically significantly non-zero (each with a p-value less than 0.05) and positive, thus implying increased probability of detection versus the standard sequence A{JB}. This means that faults in Pattern A were detected *less frequently* in the standard A{JB} sequence, when the Pattern A assembly was observed first and the JPA appeared in the second trial on Pattern B. This suggests that the JPA may have had only a limited effect in this experiment. Further, the values of the three  $\gamma$  parameter estimates in Table 7 are statistically indistinguishable from one another (the standard error estimates are essentially identical). Thus, the probability of detection of Pattern A faults is impacted by *both* the sequence of observation (whether or not the assembly of pattern A was observed before the assembly of pattern B) *and* whether or not a JPA was present.

The only distinguishable effect of the JPA is its presence in a *3-way interaction* between sequence, presence/absence of a JPA, and Pattern A. If there was an effect from *only* the JPA, the  $\hat{\gamma}_{B\{JA\}}$  and  $\hat{\gamma}_{\{JA\}B}$ , terms in Table 7 would be statistically significant and positive while  $\hat{\gamma}_{\{JB\}A}$  would be near zero. These results for Pattern A further suggest that the learning effect has a greater impact than the JPA. In contrast, the effects of sequence and/or presence/absence of a JPA on the probability of detection for pattern B (see Table 8) were not observed (i.e.,  $\hat{\gamma}_{B\{JA\}}$ ,  $\hat{\gamma}_{\{JA\}B}$ , and  $\hat{\gamma}_{\{JB\}A}$  were all statistically indistinguishable from zero with non-significant p-values).

As seen in Table 6, marking faults (fault type 1) were frequently missed and therefore dominate the 3-way interaction term in the results for Pattern A. This can be seen in the  $\beta_3$  term in Table 7 with the statistically significant p-value  $< 0.0005$  and a large positive value for the estimate (4.967). This indicates that for Pattern A, participants were *much* more likely to detect fault #3 (and fault type 3, which is an incorrect-order fault) compared to fault #1 (the default condition, as well as a marking fault). Similarly, the  $\beta_{10}$  term for Pattern B in Table 8 is statistically significant (p=0.002) with a large *negative* estimate (-2.150) and signifies that the marking fault #10 was much *less* likely to be detected than the incorrect-piece fault (fault type 2) in the default condition. This suggests that the reason Pattern A appears in the 3-way interaction term is because it has more marking errors (even when considering fault #12, which was designed to be noticed) and thus gives a better opportunity to detect differences in the probability of detection between the different fault types.

Given these results, the primary null hypothesis was not fully rejected:

H<sub>1</sub>: The presence of a JPA has no effect on the detection of faults in the quality assurance role.

The effect of the JPA was not statistically significant on its own but rather as part of a 3-way interaction between sequence of observation, presence of a JPA, and Pattern A. Thus, the additional null hypotheses were only partially accepted:

H<sub>2</sub>: There is no difference in the detection of faults with a JPA than without a JPA.

H<sub>3</sub>: The order of presentation of the JPA has no effect on the detection of faults in the quality assurance role.

H<sub>4</sub>: The order of presentation of the different patterns has no effect on the detection of faults in the quality assurance role.

What may be the most intriguing discovery in this study, however, is the ability to *reject* the final null hypothesis:

H<sub>5</sub>: The probability of detection for each fault is equal.

In hindsight, one weakness in the experimental design is that faults were not evenly distributed throughout the individual trial for each pattern, either by elapsed time or number of pieces of assembly. This might have been negligible if the assumption of average probability of detection across all faults was realistic; however, this assumption was shown to be demonstrably *wrong* and therefore acted as a confound influencing the expected results. Clearly, some fault types in this study were more difficult to detect than other fault types. The discovery that there are different probabilities of detection for *different types of error*, using this study's experimental design and within the verification context, may be the most useful result from this experiment.

Returning to the original motivation for the study provides a reminder that this research project is not just an experiment about Lego™ pieces and JPAs, but also a study of QA practices. In order for QA practitioners to effectively use tools to control risk they must be both well designed and well understood by the designer and the user. Stammers et al. (1990) suggests that task analysis methodologies themselves are limited by the difficulty of accurately predicting behavior due to the influence of cognitive factors that cannot be easily observed or modeled, thus even proper JPA development using task analyses face the same limitations. If this study were designed to detect only the main effect of a JPA on performance, then the results would have been disappointing and perhaps no interaction with other independent variables would have been discovered. These findings illustrate that this specific JPA format (i.e., a checklist) is not necessarily an effective control to improve the performance of concurrent dual verification activities. Past findings (Clark, 1982; Shriver et al., 1982; Smillie, 1985) echo the conclusion of this research that no single JPA format is best for all circumstances. It should also be noted that the verification techniques appearing in many guidebooks, such as the DOE *Human Performance Handbook* (DOE, 2009) mentioned earlier in Chapter 2, are presented as being effective without any qualifications as to whether they are actually useful or conditions under which they must be applied in order to be beneficial. This study found, however, that concurrent dual verification may not always be as effective a technique as expected or desired, either with or without a JPA.

## 6. CONCLUSIONS AND FUTURE RESEARCH

The research presented in this thesis is the first known example of evidence-based job performance aid (JPA) use in a quality assurance setting. This study used an assembly task and checklist format as the vehicle for the experiment, and incorporated previously recognized design principles from the literature into the configuration of the JPA. The design of the task included some simplifying assumptions regarding an average probability of detection between different types of error, and this led to the primary research hypotheses unable to be proved. Nevertheless, there were 3 important findings.

First, this study created a testing methodology sensitive enough to detect differences in the effects on performance between the 1) sequence of observation of patterns, 2) presence/absence of a JPA, and 3) Pattern A. The author reflects that if the *main effect* of a JPA on performance of a concurrent dual verification task were easily identifiable then it would likely have been detected long ago, though it is worth noting that there has not been any research located on concurrent dual verification either with or without the use of JPAs. Second, the results indicate that concurrent dual verification *itself* is not necessarily as effective a control as presented in the literature (however, some error detection is probably better than none in most cases). One cannot assume that having a QA checker in place will have the desired impact on error detection or mitigation, especially for specific types of errors. Third, the assumption of average probability of detection between different types of error may have been demonstrably wrong but the error detection probability could be empirically *verified* within the error construct created during this experiment. Because the probability of detection of different error types has not previously been studied in this context, it is hoped that these results are well-received by a wide community of practitioners. In addition, the subtle and complex interactions of JPA design, error type, and base error probability identified in this study help to clarify the lack of clear findings in past research.

These results can be used as a cautionary tale for processes and tools that QA practitioners use to prevent or mitigate human errors. Once an error is found, the simple conclusion is usually to blame the assembler (or checker) without understanding the complexity of tasks that both are being asked to perform. Caldwell (2008) reminds us to look for patterns of events from multiple causal factors, but this study contributes another dimension by recognizing that different types of error may require different mitigation techniques. The common risk matrix (Cox, 2008) describes the likelihood of errors occurring against their consequences, but this is *independent* of the probability of those errors being detected during a concurrent checking task. The likelihood of an error occurring can never be zero but if the probability of an error being *detected* can be improved, then an effective use of QA controls might be to improve the chances that catastrophic errors do not go unnoticed. The recognition in this study that different types of error have different probabilities of detection, which could then be inserted into a checklist as behavioral cues, might be highly useful as a leverage to the success of both JPA design and concurrent dual verification activities.

Both the methodology and results of this study are an effective baseline from which to launch future research activities. Results suggest that a learning effect overshadowed the presence of the JPA, therefore a redesigned study to counterbalance this effect might be useful. Repeating the same experiment with a uniform fault type, which would likely provide a more constant

probability of detection for all faults, may also yield intriguing results. A focus on marking faults, which can be considered signals, might provide an attractive opportunity. Since incorrect markings may fall below the signal detection threshold (Swets, 1964; Baker, 1975), such a focused study might extend this essential paradigm of vigilance theory (e.g., Mackie, 1964; Kibler, 1965; Stanislaw, 1990; Caldwell, 1999). If a similar experiment equally spaced the faults (possibly by elapsed time) throughout the assembly task, then perhaps a vigilance decrement could be studied in a number of ways. It was mentioned previously that fault #12 was designed to address findings by Tsao, Drury, and Morawski (1979) and a later study by Tsao and Wang (1984) examined faults of different difficulty, suggesting that these faults may have different probabilities of detection. These and other vigilance studies may benefit from the recognition of different probabilities of detection for different types of error, both within and outside of the concurrent dual verification construct. Altmann (2002) has fundamental areas of overlap with the current research by examining a performance decline within the current task as memory decays (or loses activation) to make room for encoding the next task. Similarly, Altmann and Trafton (2007) discuss the effects of delayed response and performance recovery after interruptions. While this study was unable to directly test the activation or interruption recovery concept, given that the presence of a JPA is not a repetitive intervention, a future experiment with multiple JPAs could address this activation model. Finally, it is worth noting that the QA checkers in this experiment used a simple checklist as the JPA, and future studies may have different results with other JPA formats. Perhaps as a result of such studies, optimal JPA formats for different fault types may emerge.

## 7. REFERENCES

- Agresti, A. (2013). *Categorical Data Analysis* (3<sup>rd</sup> ed.). New York: John Wiley.
- Allen Jr, J.P., & Rankin, W.L. (1997). Use of the Maintenance Error Decision Aid (MEDA) to Enhance Safety and Reliability and Reduce Costs in the Commercial Aviation industry. In *The 10th Annual FAA Conference on Aviation Maintenance and Inspection*.
- Altmann, E.M. (2002). Functional Decay of Memory for Tasks. *Psychological Research*, 66(4), 287-297.
- Altmann, E.M., & Trafton, J.G. (2007). Timecourse of recovery from task interruption: Data and a model. *Psychonomic Bulletin & Review*, 14(6), 1079-1084.
- Baddeley, A.D., & Colquhoun, W.P. (1969). Signal Probability and Vigilance: A Reappraisal of the 'Signal-Rate' Effect. *British Journal of Psychology*, 60(2), 169-178.
- Baker, E.M. (1975). Signal Detection Theory Analysis of Quality Control Inspector Performance. *Journal of Quality Technology*, 7, 62-71.
- Barshi, I., & Healy, A.F. (1993). Checklist Procedures and the Cost of Automaticity. *Memory & Cognition*, 21(4), 496-505.
- Booher, H.R. (1975). Relative Comprehensibility of Pictorial Information and Printed Words in Proceduralized Instructions. *Human Factors*, 17(3), 266-277.
- Booher, H.R. (1978). *Job Performance Aids: Research and Technology State-of-the-Art*. Technical Report NPRDC-TR-78-26. San Diego, CA: Navy Personnel Research and Development Center.
- Caldwell, B.S. (1999). Dynamics of Supervisory Control Task Performance: SCAMPI Project Summary. In *Automation Technology and Human Performance: Current Research and Trends; Proceedings of the Third Conference on Automation Technology and Human Performance*, Norfolk, VA, March 25-28, 1998, 213-217.
- Caldwell, B.S. (2008). Tools for Developing a Quality Management Program: Human Factors and Systems Engineering Tools. *International Journal of Radiation Oncology \* Biology \* Physics*, 71(1), S191-S194.
- Chapanis, A. (1970). Systems Staffing. In K.B. DeGreene (Ed.). *Systems Psychology*, 357-382. New York: McGraw-Hill.
- Cheng, P.G., Hecht, T.C., & Bohner, J.J. (2009) *Why Satellites Fail: Lessons for Mission Success*. Report TOR-2009(8617)-8704. El Segundo, CA: The Aerospace Corporation.
- Clark, J.E. (1982). *JPAs: An Actual Experience*. No. UNI-SA-96; CONF-820223-1. Richland, WA: UNC Nuclear Industries, Inc.
- Cox, L.A., Jr. (2008). What's Wrong With Risk Matrices? *Risk Analysis*, 28(2), 497-512.
- Davies, I.K. (1973). Task Analysis. *AV Communication Review*, 21(1), 73-86.
- Dhillon, B.S. (2008). *Reliability Technology, Human Error, and Quality in Health Care*. CRC Press.

- DOE. (1993). *Guide to Good Practices for Independent Verification*. DOE-STD-1036-93. Washington, DC: United States Department of Energy Technical Standards Program.
- DOE. (2009). *Human Performance Improvement Handbook, Volume 2: Human Performance Tools for Individuals, Work Teams and Management*. DOE-HDBK-1028-2009. Washington, DC: United States Department of Energy Technical Standards Program.
- Drury, C.G. (1982). Improving Inspection Performance. In G. Salvendy (Ed.), *Handbook of Industrial Engineering*. New York: John Wiley.
- Elliott, E. (1960). Perception and Alertness. *Ergonomics*, 3(4), 357-364.
- Elliott, T.K., and Joyce, R.P. (1971) An Experimental Evaluation of a Method for Simplifying Electronic Maintenance. *Human Factors*, 13(3), 217-227.
- Foley, J.P., Jr. (1972). *Description and Results of the Air Force Research and Development Program for the Improvement of Maintenance Efficiency*. Technical Report AFHRL-TR-72-72. Wright-Patterson AFB, OH: Advanced Systems Division, Air Force Human Resources Laboratory.
- Folley, J.D., Jr. & Shettel, H.H. (1962) *Tryout of a Preliminary Procedure for Systematically Designing Performance Aids*. Technical Report MRL-TDR62-20. Wright-Patterson AFB, OH: Behavioral Sciences Laboratory.
- Folley, J.D., Jr. (1961). *A Preliminary Procedure for Systematically Designing Job Performance Aids*. Technical Report 61-550. Wright-Patterson AFB, OH: Aeronautics Systems Division, Air Force Systems Command.
- Folley, J.D., Jr., Joyce, R.P., Mallory, W.L., & Thomas, D.L. (1971). *Fully Proceduralized Job Performance Aids: Volume 1. Draft Specification for Organizational Maintenance*. Technical Report AFHRL-TR-71-53. Wright-Patterson AFB, OH: Advanced Systems Division, Air Force Human Resources Laboratory.
- Fox, J.G. (1973). Recent Human Factors Contribution to Enhancing Industrial Quality Control. *Behaviorometric*, 3(2), 99-118.
- Gagne, R.M. (1962). Human Functions in Systems. In R.M. Gagne (Ed.) *Psychological Principles in System Development*, 35-73. New York: Holt, Rinehart, and Winston.
- Gawande, A. (2010). *The Checklist Manifesto: How to Get Things Right*. New York: Picador.
- Gebhard, R.M. (1970). *Development of a Training Program and Job Aids for Maintenance of Electronic Communication Equipment*. Technical Report HUMRRO-TR-70-19. Alexandria, VA: Human Resources Research Organization.
- Goff, J., Schlesinger, R., & Parlog, J. (1969). *Project PIMO Final Report. Volume 2. PIMO Test Summary*. Chatsworth, CA: Serendipity, Inc.
- Guttman, L. (1944). A Basis for Scaling Qualitative Data. *American Sociological Review*, 9(2), 139-150.
- Hallbert, B.P., Rodriguez, M.A., and Harbour, G.L. (1992). *Development of Job Performance Aids to Increase Human Performance Reliability: A Case Study in the Evaluation of Human Factors Principles*. Report RFP-4496; CONF-921048--1. Golden, CO: EG&G Rocky Flats, Inc.

- Harris, Douglas H. (1969). The Nature of Industrial Inspection. *Human Factors*, 11(2), 139-148.
- Haynes, A.B., Weiser, T.G., Berry, W.R., Lipsitz, S.R., Breizat, A.H.S., Dellinger, E.P., Herbosa, T., Joseph, S., Kibatala, P.L., Lapitan, M.C.M., Merry, A.F., Moorthy, K., Reznick, R.K., Taylor, B., & Gawande, A.A. (2009). A Surgical Safety Checklist to Reduce Morbidity and Mortality in a Global Population. *New England Journal of Medicine*, 360(5), 491-499.
- Helmreich, R.L. (2000). On Error Management: Lessons From Aviation. *British Medical Journal*, 320(7237), 781-785.
- Helmreich, R.L., & Merritt, A.C. (2001). *Culture at Work in Aviation and Medicine: National, Organizational and Professional Influences*. Farnham, UK: Ashgate Publishing.
- Helmreich, R.L., & Schaefer, H.G. (1994). Team Performance in the Operating Room. In M.S. Bogner (Ed.) *Human Error in Medicine*. Hillside, NJ: Lawrence Erlbaum.
- Herzberg, F.I. (1966). *Work and the Nature of Man*. New York: World Publishing.
- Joyce, R.P., Chenzoff, A.P., Mulligan, J.F., & Mallory, W.J. (1973a). *Fully Proceduralized Job Performance Aids: Vol. I. Draft Military Specifications for Organizational Level and Intermediate Maintenance*. Technical Report AFHRL-TR-73-43(I). Wright-Patterson AFB, OH: Advanced Systems Division, Air Force Human Resources Laboratory.
- Joyce, R.P., Chenzoff, A.P., Mulligan, J.F., & Mallory, W.J. (1973b). *Fully Proceduralized Job Performance Aids: Vol. II. Handbook for JPA Developers*. Technical Report AFHRL-TR-73-43(II). Wright-Patterson AFB, OH: Advanced Systems Division, Air Force Human Resources Laboratory.
- Kemeny, J.G. (1979). *The Need for Change, the Legacy of TMI: Report of the President's Commission on the Accident at Three Mile Island*. Washington, DC: U.S. Government Printing Office.
- Kibler, A.W. (1965). The Relevance of Vigilance Research to Aerospace Monitoring Tasks. *Human Factors*, 7(2), 93-99.
- Lucaccini, L F., Freedy, A., & Lyman, J. (1968). Motivational Factors in Vigilance: Effects of Instructions on Performance in a Complex Vigilance Task. *Perceptual and Motor Skills*, 26(3), 783-786.
- Mackie, R.R. (1964). *Six Years of Research on Human Factor Problems in ASW: A Summary*. Technical Report 206-28. Los Angeles, CA: Human Factors Research Inc.
- Mackworth, N.H (1950). Researches on the Measurement of Human Performance. *Medical Research Council Special Report Series*, No. 268.
- Mager, R.F. (1962). *Preparing Objectives for Programmed Instruction*. Belmont, CA: Fearon Publishers, Inc.
- Mattson, R.J. *An Assessment of Criticality Safety at the Department of Energy Rocky Flats Plant*. SCIE-DOE-201-89, July-September, 1989.
- McGrath, J.J., & Harabedian, A. (1961). Signal Detection as Function of Intersignal Interval Duration. In *Human Factors Problems in Antisubmarine Warfare* (Technical Report #9). Los Angeles, CA: Human Factors Research, Inc.
- McKenzie, R.M. (1958). On the Accuracy of Inspectors. *Ergonomics* 1(3), 258-272.

- Meilinger, P.S. (2004). When the Fortress Went Down. *Air Force Magazine*, 87, 78-82.
- Miller, R.B. (1953). *A Method for Man-Machine Task Analysis*. Technical Report 53-137. Wright-Patterson AFB, OH: Wright Air Development Center.
- Miller, R.B. (1956). *A Suggested Guide to the Preparation of Handbooks of Job Instructions*. Pittsburgh, PA: American Institute for Research.
- Murray, S.A., & Caldwell, B.S. (1996). Human Performance and Control of Multiple Systems. *Human Factors*, 38(2), 323-329.
- Newman, S.E. (1957). *On Identifying Effective Handbook Techniques* (ML-TM-57-26). Lowry Air Force Base, CO: Maintenance Laboratory, Air Force Personnel and Training Research Center.
- Norman, D.A. (1981). Categorization of Action Slips. *Psychological Review*, 88(1), 115.
- Norman, D.A. (1983). Design Rules Based on Analyses of Human Error. *Communications of the ACM*, 26(4), 254-258.
- Oberhettinger, D. (2012). Assuring that Lessons Learned Critical to Mission Success Get Used. In *2012 IEEE Aerospace Conference*, March 4-9, Paper #1177. Institute of Electrical and Electronics Engineers.
- Official: Genesis Pre-Launch Test Skipped. (2006, January 7) *Space.com* Retrieved from [http://web.archive.org/web/20060110200423/http://www.space.com/missionlaunches/ap\\_060107\\_genesis\\_update.html](http://web.archive.org/web/20060110200423/http://www.space.com/missionlaunches/ap_060107_genesis_update.html)
- Posner, M.I. (1964). Information Reduction in the Analysis of Sequential Tasks. *Psychological Review*, 71(6), 491-504.
- Rasch, G. (1960). *Studies in Mathematical Psychology: Probabilistic Models for Some Intelligence and Attainment Tests*. Oxford, England: Nielsen & Lydiche.
- Rasmussen, J. (1982). Human Errors: A Taxonomy for Describing Human Malfunction in Industrial Installations. *Journal of Occupational Accidents*, 4(2), 311-333.
- Rasmussen, J., & Vicente, K.J. (1989). Coping with Human Errors Through System Design: Implications for Ecological Interface Design. *International Journal of Man-Machine Studies*, 31(5), 517-534.
- Reason, J. (1990). *Human Error*. London, UK: Cambridge University Press.
- Russians Miswire Satellite. (1996, March 21) *The Moscow Times*. Retrieved from <http://www.themoscowtimes.com/news/article/russians-miswire-satellite/326829.html>
- See, J.E. (2012) *Visual Inspection: A Review of the Literature*. SAND2012-8590. Albuquerque, NM: Sandia National Laboratories.
- Shriver, E.L., Fink, C.D., & Trexler, R.C. (1959). *Increasing Electronics Maintenance Proficiency Through Cue-Response Analysis*. Alexandria, VA: Human Resources Research Organization.
- Shriver, E.L. (1960). *Determining Training Requirements for Electronic System Maintenance: Development and Test of a New Method of Skill and Knowledge Analysis*. Technical Report 63. Alexandria, VA: Human Resources Research Organization.

- Shriver, E.L. (1975). *Fully Proceduralized Job Performance Aids: Guidance for Performing Behavioral Analyses of Tasks*. Falls Church, VA: URS/Matrix Research Company.
- Shriver, E.L., Zach, S.E., and Foley Jr, J.P. (1982). *Test of Job Performance Aids for Power Plants*. No. EPRI-NP-2676. Alexandria, VA: Kinton, Inc.
- Smillie, R.J. (1978). Continuing Research in Job Performance Aids: The Interaction of Speed Stress and Media. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 22(1), 502-506. SAGE Publications.
- Smillie, R.J. (1985). Design Strategies for Job Performance Aids. In T.M. Duffy & R.W. Waller (Eds.) *Designing Usable Texts*, 213-241. Academic Press, Inc.
- Smillie, R.J., & Ayoub, M.A. (1980). Job Performance Aids: Evaluation of Design Alternatives Via Network Simulation. *Ergonomics*, 23(4), 319-339.
- Stammers, R.B., Carey, M.S., & Astley, J.A. (1990). Task Analysis. In: J.R. Wilson & E.N. Corlen (Eds.) *Evaluation of the Human Work. A practical Ergonomics Methodology*, 134-160. London, UK: Taylor & Francis Ltd.
- Stanislaw, Harold. (1995). Effect of Type of Task and Number of Inspectors on Performance of an Industrial Inspection-type Task. *Human Factors*, 37(1), 182-192.
- Stevens-Adams, S., Cole, K., Wenner, C., & Dornburg, C. (2013). *The Effectiveness of Verification in Industrial Applications*. Unpublished manuscript, Sandia National Laboratories.
- Swets, J.A. (1964). *Signal Detection and Recognition by Human Observers*, New York: Wiley & Sons.
- Thomson, R.M. (1972). Design of Multi-Man Machine Work Areas. In H.P. Van Cott and R.G. Kinkade (Eds.) *Human Engineering Guide To Equipment Design* (Chapter 10). Washington, DC: U.S. Government Printing Office.
- Tsao, Y.C., & Wang, T.G. (1984). Inspectors' Stopping Policy After Fault Detection. *Human Factors*, 26(6), 649-657.
- Tsao, Y.C., Drury, C.G., & Morawski, T.B. (1979). Human Performance in Sampling Inspection. *Human Factors*, 21(1), 99-105.
- White, R.E., Trbovich, P.L., Easty, A.C., Savage, P., Trip, K., & Hyland, S. (2010). Checking it Twice. *Quality and Safety in Health Care*, 19(6), 562-567.
- Wiener, E.L. (1975). Individual and Group Differences in Inspection. In C.G. Drury & J.G. Fox (Eds.) *Human Reliability in Quality Control*, 101-122. Halsted Press.

This Page Intentionally Left Blank

## APPENDIX A. JOB PERFORMANCE AID

# **Job Performance Aid**

## **For the Quality Assurance Role**

- Your role as an observer is an essential part of this important task. Complex assemblies require a second set of eyes in order to catch any errors.
- Pay attention for the following types of error:
  - An incorrect piece is installed, meaning that it is either the wrong size, wrong color, or wrong markings
  - The correct piece is installed, but in the wrong orientation
  - The correct piece is installed, but in the wrong location
- Feel free to ask questions about the task at any time. If necessary, ask the assembler to stop until you are comfortable with proceeding.
- The assembler should not turn to the next page of the instructions without your approval.
- For each page of the instructions, the order of assembly does not matter.
- The box contains 512 total parts. Some parts will be used and some will not.

This Page Intentionally Left Blank

# APPENDIX B. FAULT TYPES

Fault #1, detailed in Chapter 3, is reproduced here:

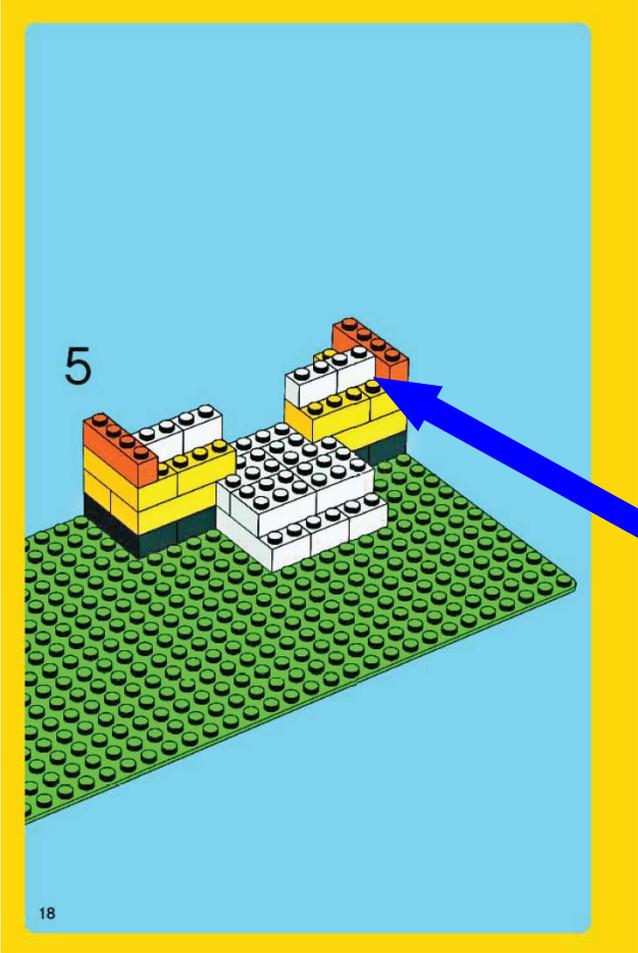
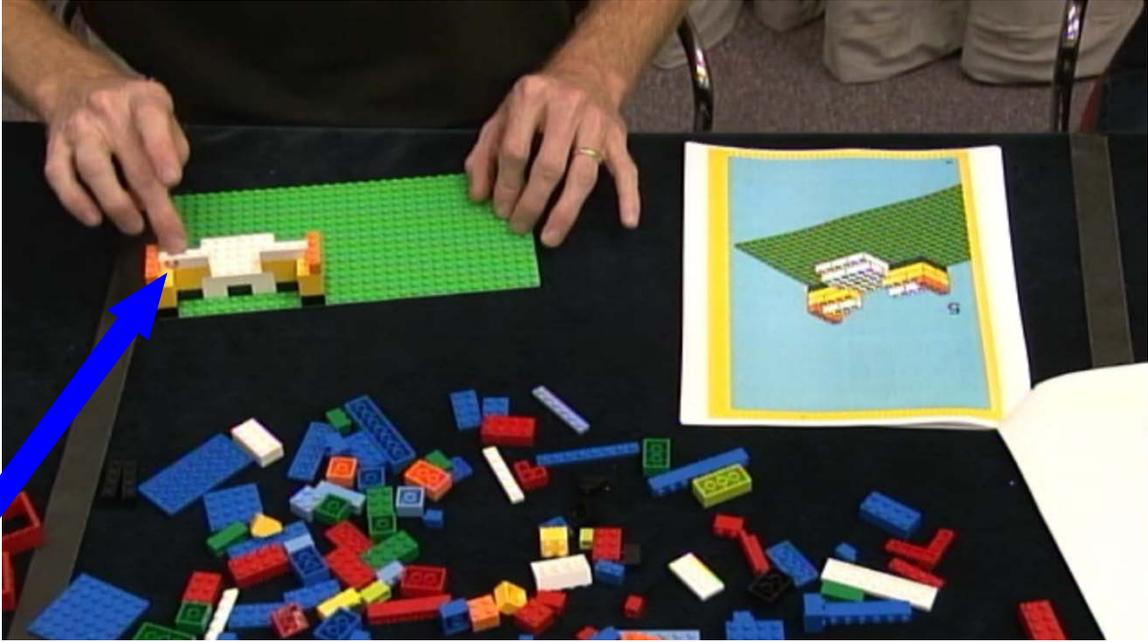


Figure B1: Sample Instructions Page, Depicting Fault #1 (Fault Type 1)



**Figure B2: Fault #1 Occurring in an Experimental Trial**

Figure B3 shows a page from the assembly instructions for fault type 2 (Fault #8), which is where an incorrect assembly piece is used. The arrow points to the black 2 x 4 piece that is replaced by two 1 x 4 pieces in the experimental trial, one of which is shown in Figure B4.

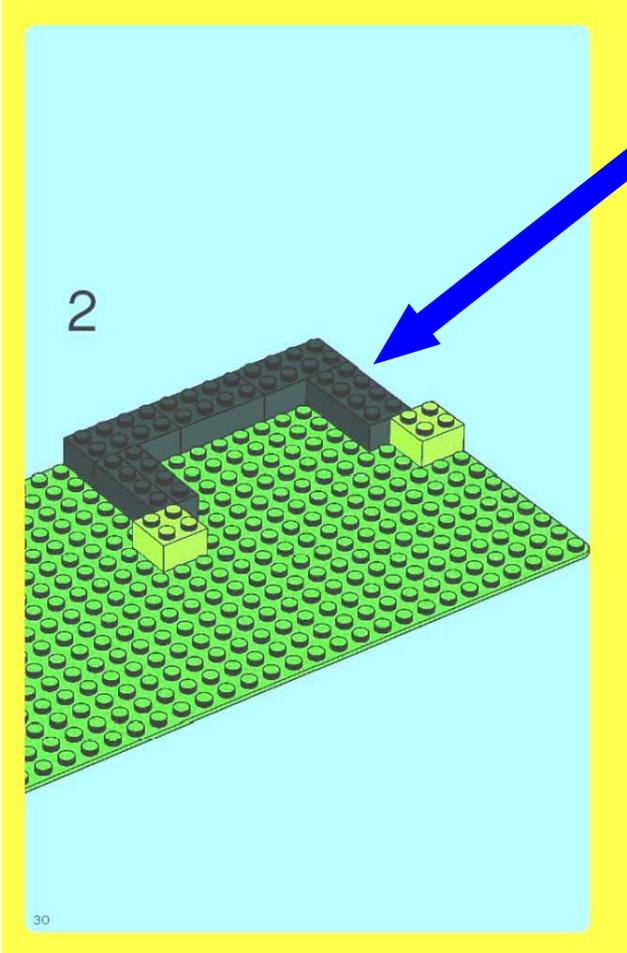
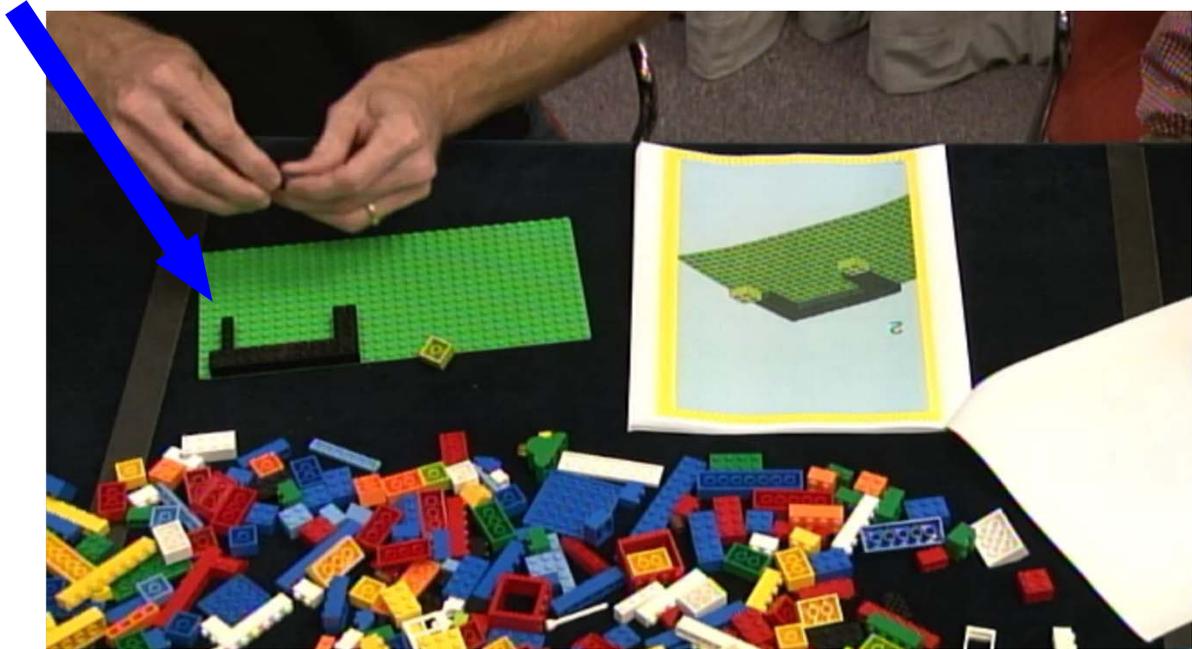
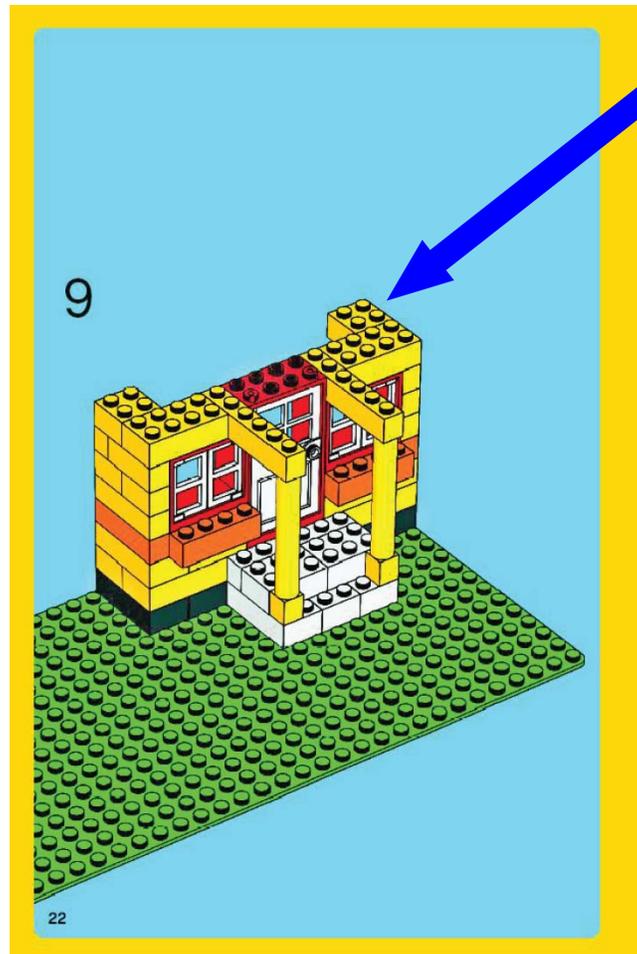


Figure B3: Sample Instructions Page, Depicting Fault #8 (Fault Type 2)



**Figure B4: Fault #8 Occurring in an Experimental Trial**

Figure B5 depicts fault type 3 (also Fault #3) whereby the correct pieces are used but they are not in the correct order. The arrow shows the correct orientation where the 2 x 2 piece is in the rear corner of the assembly. Figure B6, however, indicates the fault whereby the 2 x 4 piece is moved such that it appears in the rear corner instead.



**Figure B5: Sample Instructions Page, Depicting Fault #3 (Fault Type 3)**



**Figure B6: Fault #3 Occurring in an Experimental Trial**

Figure B7 shows a red window in the assembly that is correctly installed. An example of fault type 4, which uses a correct part but in the wrong orientation, is seen in Figure B8 where the window is installed backwards.

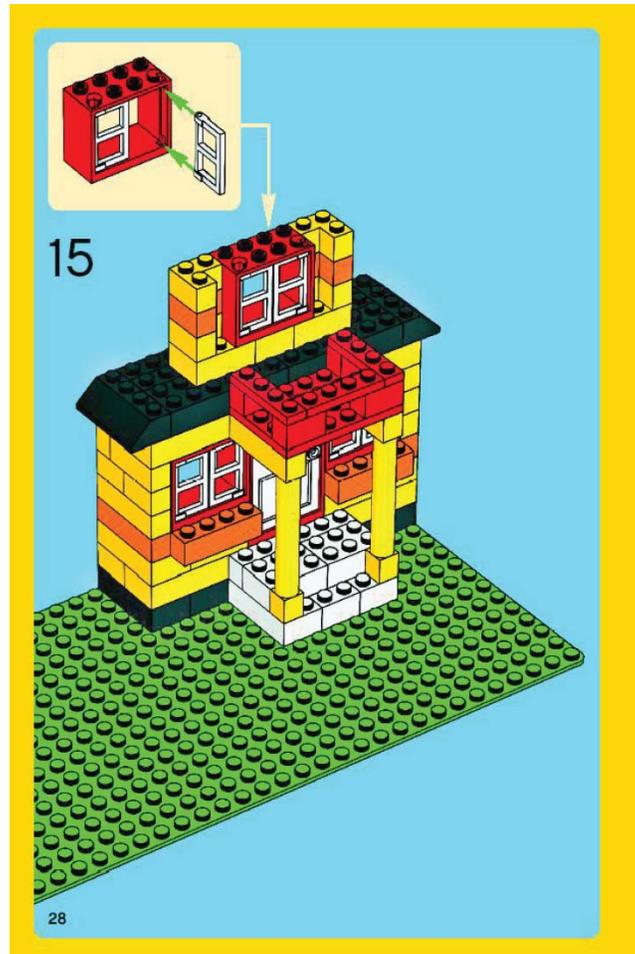


Figure B7: Sample Instructions Page, Depicting Fault #5 (Fault Type 4)



**Figure B8: Fault #5 Occurring in an Experimental Trial**

## DISTRIBUTION

3 Purdue University, School of Industrial Engineering  
Attn: Barrett Caldwell (1)  
Mark Lehto (1)  
Archive Copy (1)  
315 N. Grant Street  
West Lafayette, IN 47907-2023

1	MS0127	Caren Wenner	0254
1	MS0330	Brent Blankenship	0427
1	MS0330	Pete Geib	0422
1	MS0330	Ron Farmer	0420
1	MS0330	Maria Galavíz	0426
1	MS0330	Marcey Hoover	0420
1	MS0330	Ted Kreifels	0424
1	MS0330	Mike Lopez	0421
1	MS0330	Ruben Muniz	0425
1	MS0330	Nayeli Perez	0428
1	MS0330	Leslie Simmons	0421
1	MS0330	Jeffrey West	0421
1	MS0428	Jennifer Lindblom	0428
1	MS0428	Jennifer Samora	0428
1	MS0748	Susan Stevens-Adams	6231
1	MS0829	Ed Thomas	0415
1	MS0830	Kerstan Cole	0431
1	MS0830	Courtney Dornburg	0431
1	MS0830	Judi See	0431
1	MS0872	Dennis Owens	0423
1	MS0899	Technical Library	9536 (electronic copy)



**Sandia National Laboratories**