Assessing Training: A Methodology for Analyzing Tacit Assessment

Phillip N. Jones, Thomas Mastaglio, PhD
MYMIC LLC
Portsmouth, Virginia
phillip.jones@mymic.net, tom.mastaglio@mymic.net

ABSTRACT

This paper describes assessing the effectiveness of training programs using the Study of Organization Opinion (SO3) methodology which has been developed in the course of performing multiple assessment projects for the Army and joint community. SO3 is based on concepts that are foundational to Customer Relationship Management (CRM). SO3 is effective in extracting tacit assessment knowledge from training stakeholders and analyzing the knowledge to provide actionable information to decision makers.

The SO3 methodology uses open-population surveys of training stakeholders. Organizational objectives (study goals) are disaggregated to derive deliverable questions and the likely stakeholder population is studied to choose respondent groups. Questions are linked to respondent groups to create questionnaires. Questionnaires are delivered via multiple means, including via the web, interview, focus groups, and written surveys. Once questionnaires are complete, the data is aggregated using various analysis methods to provide various levels of findings.

The paper provides numerous practitioner lessons learned, such as: the importance to the assessor of knowing the targeted organization, the importance of treating assessment as a cross-organizational competency, and the emphases on response quality. Finally, the paper recommends assessment improvement objectives, such as seamlessly integrating assessment into the training program, persistent assessment, and inter training-modality metrics for measuring training competency required and supported.

ABOUT THE AUTHORS

Phillip N. Jones is the Manager of the Training & Assessments laboratory within MYMIC LLC, a high technology start up specializing in modeling, simulation and visualization solutions to support training, military analysis, experimentation and business decision-support. Mr. Jones has over twenty-two years of professional experience as an Army combat-arms officer. He has extensive Army and Joint operational and training experience. This has included training Active and Reserve Component units, training from platoon to the national level, and training Dutch Marines en rout to Cambodia and serving as the advisor to the Kuwaiti Armed Forces J3. More recently, Mr. Jones has led a series of studies on training effectiveness in Army and joint units, measuring the effectiveness of training programs, identifying training support requirements in the face of tactical and strategic changes, measuring the impact of the Global War on Terror on US training and developing innovative methods for measuring training efficacy.

Dr. Thomas W. Mastaglio is the President of MYMIC LLC, a high technology start up specializing in modeling, simulation and visualization solutions to support training, military analysis, experimentation and business decision-support. Dr. Mastaglio was the Executive Director of the Virginia Modeling, Analysis and Simulation Center from November 1996 to June 2000. He is a 1969 graduate of the U.S. Military Academy, retired from the Army in 1991 then worked in the defense industry as a senior engineer, program manager, and scientist. Dr. Mastaglio earned a Doctor of Philosophy degree from the University of Colorado in Computer and Cognitive Science in 1990. His research interests include the application of artificial intelligence to improving human-computer interaction and learning, usability engineering, cognitive modeling, the development of large-scale enterprise models and

simulations, and educational requirements for modeling and simulation. Dr. Mastaglio has served as a member of the Army Science Board, a Consultant to TRADOC, a technology consultant to the Department of Defense DDR&E and as an Advisor to the Commonwealth of Virginia General Assembly's Joint Commission On Technology and Science. Dr. Mastaglio is the author or co author of over 40 papers published in professional journals and conference proceedings.

Assessing Training: A Methodology for Analyzing Tacit Assessment

Phillip N. Jones, Thomas Mastaglio, PhD
MYMIC LLC
Portsmouth, Virginia
phillip.jones@mymic.net, tom.mastaglio@mymic.net

INTRODUCTION

The training management process, as exampled by the Department of Defense's (DoD's) Joint Training System (JTS), consists of four phases, all essential to a successful training effort. These phases are: Requirements, Planning, Execution, and Assessment. Of all these phases, Assessment is perhaps the most problematic. The JTS defines Assessment as (CJCSM 3500.03A, 2002):

"A training assessment is the analytical process used by commanders to determine an organization's proficiency to accomplish the capability requirements defined in JMETs (Joint Mission Essential Tasks). The assessment phase of the JTS provides the commanders and staffs at each level of command valuable information that describes a direct payoff in terms of improved mission capability for the effort associated with the first three phases of the JTS... The goal of the assessment phase is to provide a clear structure to institutionally capture those insights and create a learning organization."

There is a concern with the level of assessment being done within the DoD. This paper postulates, however, that there is extensive training assessment occurring, at all levels and as part of every training event. This assessment, however, is tacit. It occurs within the minds of DoD uniformed and non-uniformed professionals. The challenge is accessing this tacit assessment in such a way that supports analysis and the ultimate information requirements of decision makers.

A definition for tacit knowledge is (Dictionary of Philosophy of Mind, accessed 2007): "Knowledge that enters into the production of behaviors and/or the constitution of mental states but is not ordinarily accessible to consciousness." Tacit knowledge exists within the individual, though some have postulated that it can exist within social structures. Under some organizational knowledge creation theories, all learning begins with tacit knowledge. However, it is difficult to transform tacit knowledge into explicit knowledge (Polyani, 1966 and Nonaka & Takeuchi, 1995). The Study of Organizational Opinion (SO3) process is a controlled tacit knowledge extraction method that is efficient, complete and effective in supporting analysis. SO3 has been applied in numerous studies of training, both for front end requirements analysis and in training program effectiveness analysis.

SO3 is the result of a development process spanning several studies. It was initiated in 2004 in response to a requirement to assess the effectiveness of the Close Combat Tactical Trainer (CCTT). An efficient method was required that could quantify the effectiveness of CCTT without expensive and disruptive experiments. The study team recommended leveraging Customer Relations Management (CRM) techniques. These are techniques used by business to assess their effectiveness in servicing their customers. They are based on analyzed customer surveys. For applying CRM to the CCTT effectiveness problem, the team analogizing small unit leadership—platoon through battalion—as CCTT customers and surveying them. Table 1 shows the developmental history of SO3.

Table 1. SO3 Development History

2004:	Assessing the Effectiveness of the Close Combat Tactical Trainer
	Soldier CATT Front End Analysis
2005:	Contribution of Virtual Simulation to Combat Effectiveness--Experimented with web delivery
	Contribution of Virtual Simulation to Combat Effectiveness within the Army National Guard--Integrated web delivery
2006:	Evaluating the Close Combat Tactical Trainer (CCTT) Staffing and Redistribution
	Current Operational Environment Impact Upon The Soldier Combined Arms Tactical Trainer
	Joint Training Assessment Study--SIPRnet delivery, operational level assessment
2007:	Emergency Management Collaborative Planning Tool--Non-DOD use
	Software to Assess Readiness and Train Medical Support Operations Teams
	Consultation to Eastern Virginia Medical School Medical Simulation Program

Training Assessment Insights

Assessing training effectiveness means providing training stakeholders with information sufficient to make appropriate decisions. Assessment must support the spectrum of training stakeholders, including: trainees, trainers, training managers, training developers and training theorists.

Another role of training assessment is to assist in identifying the Return on Investment (ROI) or Return on Expectation (ROE) of training. "Return" is the value of training received from an event or process, which assessment should be able to evaluate. "Investment" means the total resources placed into a training modality. "Expectation" is the training value anticipated by training stakeholders, and can be assessed against Return.

SO3 METHODOLOGY

The SO3 process has been developed to ensure a complete and thorough assessment. SO3 consists of seven steps divided into three phases: Preparation, Execution, and Analysis & Reporting. These steps are described below:

Figure 1. SO3 Methodology

Preparation
- **Step 1**: Map the organization and organizational goals
- **Step 2**: Map the respondents
- **Step 3**: Disaggregate goals to questions
- **Step 4**: Link questions to respondent to develop questionnaires

Execution
- **Step 5**: Deploy, track, and retrieve questionnaires

Analysis & Reporting
- **Step 6**: Aggregate and analyze data points/answers
- **Step 7**: Report on results

Step 1: Map the Organization

Mapping the organization is the most important and the most difficult step. The assessment team must start with organization objectives, especially change goals. From there, the assessment must work hard to understand everything possible about the organization, keeping in mind that the assessment may have to assist the organization in understanding or even identifying its goals. Some organizational aspects to consider during mapping include:

- What are the organizational goals? What do they want to change?
- What is the organizational structure? Who makes decisions? What are those decisions? What information do they need?
- What are the existing assessment/decision processes? How important is continuity versus evolution? What management processes must the information fit into?
- What are the various organizational agendas? How do you work around them?

It is important to realize that in any assessment opportunity, there is probably already existing management processes, such as Lean, Lean Six Sigma, Balanced Scorecard, or others. The assessment team must discover, understand, and synchronize itself with these existing management processes.

The team must also be aware of pre-existing assessment efforts at the organization. Continuity of assessment will be important to the organization and there will be resistance against new assessment methodologies. It should be understood, though, that continuity lies in the information and insight which comes from the assessment, not in the particular assessment methodology, which is just a means to the end. Continuity of information can be maintained without rejecting newer and more effective assessment methodologies.

Step 2: Map the Respondent

Concurrent with mapping the organization, the assessment team must examine the population which contains the desired tacit knowledge and which will be the focus of the assessment collection. The team must define this population, focusing on those individuals who are likely to have the desired information and who the assessment effort can reasonably reach. The team works with the organization to identify those demographic factors—age, gender, rank, specialty, etc—that are important to understanding population responses. SO3 uses relative responses between these factors as a primary means of analysis. Finally, the team identifies the most significant or critical demographic factor and divides the population based on this factor into different Respondent Groups. In past work, Respondent Groups were divided by duty position, rank, or specialty.

Step 3: Disaggregate Goals into Questions

Organizational assessment goals are usually broad, high-level questions. These questions tend to be so broad that they cannot reasonably be delivered to an individual and elicit specific, analyzable data. Rather, the assessment team must disaggreagate these broad goals through sufficient levels that deliverable questions—questions that respondents could be expected to understand and answer with usable data—are identified. The assessment normally disaggregates goals through three levels: Issues, Sub-Issues, and Questions—all in the form of a query. Each level supports the one above, so the questions support Sub-Issues and the Sub-issues support the Issues, which provide the insight required to fulfill the study goals and inform organizational change. Normally, each query is supported by three or four subordinate queries, though they may be supported by fewer. Too many supporting queries tends to make the process unmanageable. Figure 2 shows this disaggregation:

Figure 2: Disaggregation of Study Goals

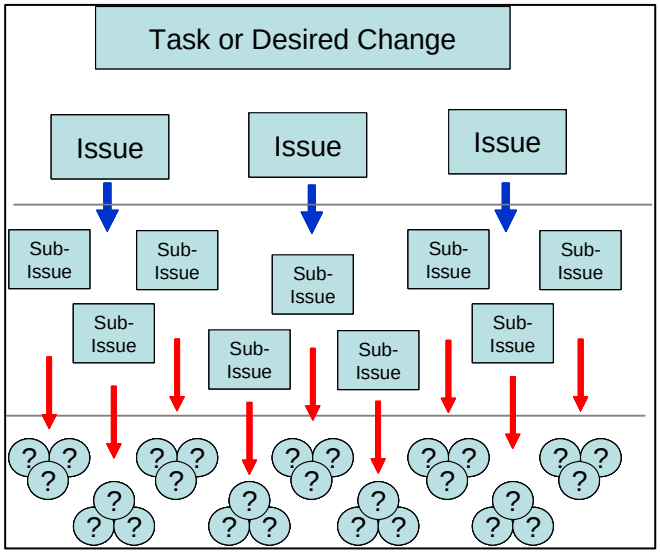

Question preparation is a critical art. Questions are driven by the information requirement. This has several impacts upon the form of Questions. First, Questions should be written for best understanding by the intended respondent. One Question may be delivered to multiple respondents, who may have significantly different perspectives and understandings. In this case, the Question should be phrased for each respondent group, keeping the information requirement constant across the differently phrased questions. As an example, if the information requirement is to determine the theater-specific realism of a training program, the question should be written in future tense for those who have not deployed to theater and present tense for those in theater.

Question form must also support collection of the required information. Question form refers to the type of question asked. There are many possible types: Yes/No, True/False, multiple choice, rank order, fill in the blank, essay, etc (Schuman & Presser, 2002 and DeVellis, 2003). Effective assessments are open to all forms of questions and may include open and closed questions, and qualifiable and quantifiable questions. Open, qualifiable questions add to the analysis burden of the assessment team, as well as to the burden of the respondent, since they are generally more difficult to answer. The assessment team must be very cognizant of the total burden upon the respondent for reasons to be discussed below. As for the burden to the analysis, that is limited somewhat by the phenomenon that by asking a question, even an open question, the asker creates a finite set of possible, valid answers (valid as defined by containing information useful to the assessment).

In studies of complex organizations, with various echelons of command or management and with various channels of authority, the disaggregation should be tied to those echelons, such that Issues are executive level information requirements, Sub-Issues are management level requirements, and Questions are targeted towards the assessment population.

Step 4: Link Questions to Respondents

The final step in preparation is to link questions to respondents to create questionnaires. Normally, more questions will be prepared than can be delivered to the assessment population, so this step requires deciding which questions NOT to deliver to a respondent group. Although it is always surprising to see how quickly a respondent can finish a questionnaire, experience has shown that a single questionnaire should not contain any more than approximately fifty questions.

A questionnaire includes demographic questions and questions developed from the disaggregation performed in Step 3. It is good practice to also include "filtering" questions. These are questions designed to inform the assessment team about biases or other critical aspects of the individual respondent. The most important filtering questions, especially when the assessment team cannot strictly control who responds to the questionnaire, are those designed to identify respondents who are not within the desired assessment population. These could

be individuals who want to positively contribute to the assessment effort, individuals who are just curious, or individuals who want to negatively contribute. These individuals are referred to as "other-respondents". Filtering questions are used to identify other-respondents, but the questions must be subtle so as not to let respondents believe they are being screened. This is done to prevent other-respondents from falsifying answers to filtering questions, and to maintain the most positive atmosphere for all respondents. The data received from other-respondents is separated from the primary data, but not necessarily eliminated. All valid data is valuable to the assessment and even data form other-respondents may contain insight.

Question sequencing on questionnaires is important (Schuman & Presser, 2002). First, experience has shown that there is a natural drop off by respondents after approximately twenty questions. This drop off can be physical—respondents departing the questionnaire—or cognitive—respondents not placing effort into their answers. For this reason, the most critical questions should be front loaded within the questionnaire. Also, answering one question will impact the answer to other questions by providing information to the respondent or biasing the respondent's subsequent answers, so questionnaires should be mindful of this and sequence questions appropriately.

Questionnaires are deployed to the respondent by several means. These include via interviews, focus groups, written surveys, and increasingly via the web. Deployment means require varying levels of resources but return varying depths of understanding. Thus, interviews are the most resource intensive but provide the most in-depth understanding of the individual tacit knowledge. The web can deliver questionnaires to thousands of respondents worldwide, but there is no means to dig into responses.

As mentioned above, the assessment team must be aware of the burden to the respondent involved in answering a questionnaire. An individual will only expend a finite amount of time and energy on any questionnaire. The assessment team should consider this finite amount as a budget from which it spends a small amount with every respondent action, both physical and mental. Thus, the team should construct its questionnaires to absolutely minimize respondent effort, such as minimizing pre-questionnaire instructions, switches from mouse to key board, etc. In addition, the assessment team should be wary of question forms which require a significant amount of mental energy, such as questions that ask the respondent to rank order several answers.

The assessment team can do things to increase the attention budget or slow its spending. Establishing a personal relationship between the respondent and the assessment effort by explaining the purpose of the effort and the importance of the respondent's information is one method of increasing the budget. Showing the respondent where he/she is in the questionnaire via a progress bar, or permitting non-linear movement through the questionnaire are methods of slowing the budget spending.

Questionnaires must be thoroughly rehearsed. This includes having exemplar respondents execute the entire questionnaire as it will be delivered to the assessment population. This allows internal feedback on the ease of the questionnaire execution and on the understanding of the questionnaire questions. In addition, each questionnaire should be delivered to an exemplar respondent via an interview, regardless of the intended deployment means. Experience has shown that the only way to truly assess a question is through verbalizing it to a respondent.

Technology allows the deployment of dynamic questionnaires. These questionnaires use respondent answers to better target specific questions to correct respondents, allowing each individual a separate path through the questionnaire or presenting a question specifically worded for the individual respondent.

Figure 3: Individual Questionnaire Paths

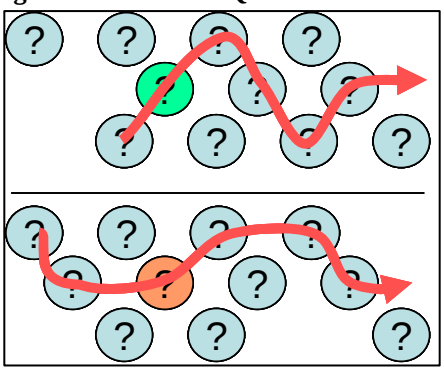

Step 5: Deploy, Track, and Receive Questionnaires

As stated above, questionnaires are deployed via interviews, focus groups, written surveys, or the web. The objective of deployment, tracking and reception is to maximize questionnaire response and throughput. "Response" is the number of respondents who come to the initiate the questionnaire. "Throughput" is the number of respondents who finish the questionnaire

providing valid data. Both metrics are important. The most important metric, however, is response quality. This is a combination of total response, response by respondent group, and the quality and quantity of data. Response quality is in contrast with response rate, which is a factor of the number or respondents and the assessment population. Response rate is extremely difficult to measure as the assessment population is open and likely to change and is dependent upon context more than any other factor (Schonlau et al, 2002). Further, response rate does not take into consideration questionnaire throughput nor data quality.

Questionnaires should be deployed through multiple methods, i.e. via multiple announcement means. This will reach the largest size of assessment population and emphasize the importance of the assessment effort. Announcements can be targeted towards a specific group of individuals, or general, such as an open announcement on a web home page. Experience shows that subsequent announcements to the same population via the same means are not highly effective, so maximum effort should be placed into the initial announcement.

As data in terms of answers come in, the assessment team should insure it is properly stored in an appropriate database with adequate backup.

Step 6: Aggregate and Analyze Data Points/Answers

Analysis within SO3 should be considered a discovery process as each action leads the analyst to a greater insight on the assessment process and the study goals. It is also the easiest part of SO3.

SO3 analysis is normally done in three sweeps. The first sweep is the preliminary sweep and is completed immediately after the cessation of data collection. This sweep consists of just the closed, quantifiable questions, including demographic questions. This sweep provides an initial assessment of the data and informs later analysis. It also allows the assessment team to provide some immediate feedback to the assessed organization.

The second sweep is the primary, planned sweep. The SO3 methodology reverses the disaggregation conducted in Step 3, aggregating the data into Consensus Answers for each question, Sub-Issue Findings, and Issue Findings as shown in figure 4.

Each individual answer for each question is examined by an analyst, who uses a Data Analysis Interface (DAI) to combine individual answers into a respondent group consensus answer. For closed-end answers, this is a matter of simple math. For open-end answers, the DAI assists the analyst in identifying the separate parts of a respondent's answer. In a single answer, a respondent might provide multiple recommendations or bits of knowledge. These are combined and added across the respondent group to provide weighted answers, showing both the emphasis and totality of the respondent group's opinion or knowledge. Note that even a low weight answer, one possibly mentioned only once by a single individual respondent, is captured and maintained within the analysis because the analyst cannot know if that single datum is not a point of inspiration.

Respondent group consensus answers are aggregated into by-respondent group Sub-Issue Findings. These are then aggregated into consensus Sub-Issue Findings, which are aggregated into Issue Findings. Multiple analysts may provide input for each aggregation, which provides differing perspectives, but a single analyst is responsible for the aggregation chain to maintain consistency and to maximize the depth of insight being gained by working the data.

Aggregation is done first by relative analysis, comparing one respondent group's consensus answer or Sub-Issue Finding to another, and then by consolidation.

The third sweep is the final inspection of the data and consists of unplanned or informal data analysis. SO3 analysis always turns up unexpected results and the third sweep follows these interesting threads. It includes of different analysis methods such as cluster analysis. The analyst also looks for alternative methods

Figure 4. Aggregation of Data

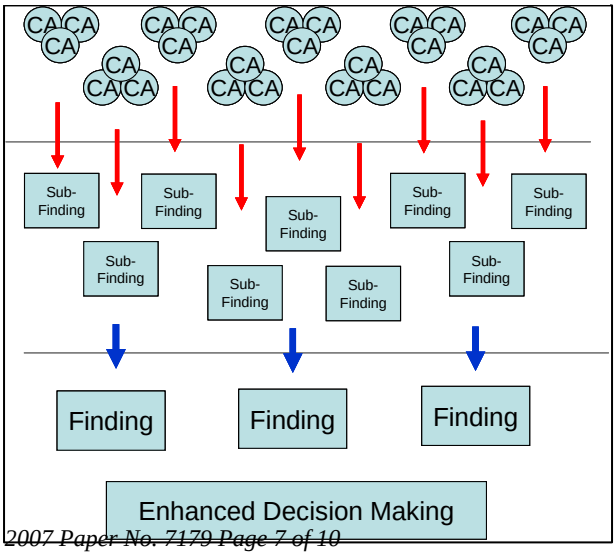

of examining the study goals, such as reviewing historical data.

Step 7: Report Results

The value of the assessment is only as great as the organization can perceive and act on it. To formulate decisions and drive change, results must be reported is a manner that supports the assessed organization's use. SO3 provides large amount of data. Even a small SO3 study can generate several hundred, to several thousand data points. Reports must be presented that maximize clarity while still allowing organizational users to drill down into the details of the data. Normally, only findings are included within the body of a report, with detailed analysis presented in appendixes.

ASSESSMENT INSIGHTS

Executing the SO3 methodology has provided some insights or recommendations into assessments, starting with internal metrics for judging the assessment effort: an assessment taxonomy.

Assessment Taxonomy

As the SO3 methodology developed over several projects, it became apparent that it required a taxonomy to serve as a metric and to demonstrate intended organizational assessment results. The below taxonomy was developed to demonstrate the differing results assessment can produce. As the levels of the assessment increase, greater insight and greater support to decision making is obtained. However, higher levels require greater resource expenditure.

Level 1: Identify
This is the simplest form of assessment. An example would be an end of course critique sheet with a single Likert scaled response. This level only provides a quick assessment without discernable, actionable information.

Level 2: Discriminate
This level represents general assessment but describes results that can be discriminated by respondent or other measure. It begins to focus the assessment to specific areas, which can be targeted for future effort.

Level 3: Explain
This level informs the organization as to why assessment results are what they are.

Level 4: Improve
An assessment at this level provides the organization with possible ways to improve the process or program being assessed.

Level 5: Transform
This is the highest level of assessment and the most difficult and resource intensive. An assessment at this level will inform the organization of ways to replace existing processes or programs with new, more effective and efficient methods.

Assessment Recommendations

In order to achieve higher level, more effective assessments, there are several improvements that should occur.

The first improvement is moving assessment to a state where it is seamlessly ingrained within the training. The SO3 methodology and other assessment methodologies tend to be discrete from the training event. They sit on the outside and look in. Further, they tend to be disruptive of the organization's normal operations. This is likely one reason for the reported difficulty in obtaining assessment mentioned at the start of this paper. Methods should be found and incorporated into the development of training modalities, of ingraining assessment such that it invisible to the training audience.

One way to do this is by continuing to take inspiration from the CRM field. The modern retail world uses the Point of Sales (POS) to collect data. Every time someone purchases an item at Wal-Mart or Sears, a small amount of data is collected. Over time, this small amount of data accumulates and starts providing significant insight to these companies (National Association of Electrical Distributors, 2005).

Every time a training audience uses a training modality, data should be collected. This data could come from the training modality. Even easier would be persistent data collection at in- and out-processing. A training manager could easily use a computer to enter the time he/she began and ended training, the number of personnel trained, and the tasks trained. He/she could also be asked a few questions. Across a training modality, this data would accumulate and an assessment could analyze it by time, location, task, or modality.

Another simple way of improving assessments is to treat them as cross-organizational competencies. Assessment should be a separate, cross-organizational

specialty. Evidence of this is found in current organizational management approaches many of which are being adopted by the military, such as Balanced Score Card or Lean Six Sigma. A central foundation of these concepts is developing the information required to most effectively achieve organizational objectives. Feedback from assessment is central in that it facilitates leaders who want to holistically examine their organizations or programs.

By treating Assessment as an independent competency, an organization will be able to acquire and apply skilled and focused assessment capability from experts in this area. This will provide high quality assessments to support program and organizational level evaluations by program managers and corporate leadership. Additionally, contracting for assessment as a singular competency will allow organizations to be agile in responding to internal and external information requirements with a result that achieves impartiality.

A final recommendation for improving assessment is the development of common terminology that links training modalities with competency requirements. With the growing number of training processes and devices, there is a requirement for a comprehensive and simple method to measure training requirements and training support, both in absolute terms and terms relative to other processes and devices. This method should have the following characteristics:

- It must generate a measurement that can serve as a common denominator across a variety of similar and dissimilar training requirements and modalities.
- It must possess sufficient rigor to establish legitimacy across the training development community.
- The measure must provide sufficient insight to support comparison between and amongst training requirements and modalities.
- The measurement should be objectively determinable by personnel with minimal instruction.

The closest the US military comes to a training scale are the Training Proficiency Evaluation (TPE) Levels. There are four levels: T—Trained, P—Partially Trained, U—Untrained and N—Not Observed. Commanders determine these levels for their units by comparing their ability to perform essential tasks against both doctrinal and mission-related performance norms.

TPE levels are not an effective means of measuring either required competence levels or levels of support to tasks. First, by their very nature, the TPE levels cannot describe the required training competency. As for measuring support of training modalities to skill competence, TPE levels are also inadequate. TPE levels are relative. They are the opinion of the commander based upon his perceptions of his organization's proficiency versus a required proficiency level derived from mission requirements. Doctrinally, TPE levels cannot be used to assess tasks separate from mission dependent requirements. Secondly, TPE does not provide sufficient information regarding the manner of support or non-support. The T-P-U provides only three gradients or proficiencies with untrained (U) covering the vast majority of proficiency levels. T-P-U does not provide sufficient granularity for program managers to adequately assess a task or process. Further, T-P-U does not indicate the reasons for the assessment sufficiently to empower process comparison or improvement.

The learning sciences have a myriad of methods for measuring learning. They can describe the nuances of learning in excellent detail. However, this level of explanation is too detailed for the comprehensive and simplistic measurement requirement described above. What is required is a unified learning scale, something conceptually similar to the unified theory being sought by physicists. What is needed is a simple scale, informed by learning science, that can be used by lay people across the training domain.

REFERENCES

DeVellis, Robert (2003) Scale Development: Theory and Applications, 2nd Ed., Sage Publications, San Diego CA

Dictionary of Philosophy of Mind – Tacit Knowledge. Retrieved 22 June 2007 from http://philosophy.uwaterloo.ca/MindDict/tacitknowledge.html

Joint Chiefs of Staff (2002). CJCSM 3500.03A Joint Training Manual for the Armed Forces of the United States. Washington, D.C., Government Printing Office

National Association of Electrical Distributors (2005) *Recommended Best Practices: Collection of Point-of-Sales (POS) or Point-of-Transfer (POT) Data*, Retrieved 22 June from http://www.naed.org/images/Admin/POSWhitepaper.pdf

Nonaka, Ikujiro & Takeuchi, Hirotaka (1995) *The Knowledge-Creating Company*, Oxford University Press, NY

Polyani, Michael (1966) *The Tacit Dimension*, Routledge & Kegan Paul, London

Schonlau, Mattias; Fricker, Ronald D. Jr. & Elliot, Marc N. (2002) *Conducting Research Surveys via E-Mail and the Web*, Rand Corporation, Santa Monica CA

Schuman, Howard & Presser, Stanley (1996) *Questions & Answers in Attitude Surveys: Experiments On Question Form, Wording, and Context*; Sage Publications, San Diego CA