

2006 CCRTS

THE STATE OF THE ART AND THE STATE OF THE PRACTICE

PAL Boot Camp: Acquiring, Training, and Deploying Systems with Learning Technology

C2 Concepts and Organizations; C2 Architecture; Policy

Douglas S. Lange

SPAWARSYSCEN San Diego

SPAWARSYSCEN

53560 Hull Street

San Diego, CA 92152-5001

619-553-6534/619-553-5322

doug.lange@navy.mil

ABSTRACT

The Defense Advanced Research Projects Agency (DARPA) has implemented a program to build the first instance of a complete cognitive agent. The program, called Personalized Assistant that Learns (PAL), is expected to yield new cognitive technology of significant value to the military. Like any good assistant, PAL must learn by observing its human master and by accepting explicit advice and instruction.

With traditional engineering projects evaluation can be done in a straightforward manner determining if the documented requirements of the system have been met. Agent-based capabilities and other network centric capabilities complicate matters because the environment that they will operate under constantly changes. Add to that complication, the ability to learn new capabilities, and testing whether or not a new agent is ready to be deployed becomes a problem beyond the current state of art and practice.

This paper lays out the problem in such a way as to identify the key issues for evaluation, transition, and acquisition. By doing so, research can be targeted for the problem and solutions found. An initial experiment design is proposed as well to examine the role that evaluation will play towards transitioning cognitive systems that learn into the military environment.

INTRODUCTION

A PERSONALIZED ASSISTANT THAT LEARNS

A cognitive system is one that can reason, using substantial amounts of appropriately represented knowledge; can learn from its experience so that it performs better tomorrow than it did today; can explain itself and be told what to do; can be aware of its own capabilities and reflect on its own behavior; and can respond robustly to surprise [Gunning 2004].

The Defense Advanced Research Projects Agency (DARPA) has implemented a program to build the first instance of a complete cognitive agent. The program, called Personalized Assistant that Learns (PAL), is expected to yield new cognitive technology of significant value not only to the military, but also to business and academic sectors. It will spur pioneering research in cognitive information processing, including areas of artificial intelligence, machine learning, knowledge representation and reasoning, machine perception, natural language processing, and human-computer interaction.

Through the PAL program, researchers will develop software that will function as an enduring personalized cognitive assistant to help decision-makers manage their world of multiple simultaneous tasks and unexpected events. PAL has two concurrent efforts underway. Carnegie Mellon University's effort under PAL is called RADAR, for Reflective Agents with Distributed Adaptive Reasoning. The system will help busy managers to cope with time-consuming tasks such as organizing their E-mail, planning meetings, allocating scarce resources such as office space, maintaining a web site, and writing quarterly reports. Like any good assistant, RADAR must learn by observing its human master and by accepting explicit advice and instruction.

SRI International is developing a cognitive assistant called CALO (Cognitive Agent that Learns and Organizes) that supports users in carrying out their routine tasks, assisting them when the unexpected happens. CALO knows things and does things. It will learn by working with, observing, and being advised by its users. In the early years, It will carry out specified tasks composed of primitive actions (e.g., receiving messages, reading messages, saving messages in folders, etc.) based on learning user preferences and taking user advice. In the later years, CALO will be more collaborative, working closely with the user to elaborate and define tasks and responses to events. CALO can also learn new ways of accomplishing objectives and will assume greater responsibility in initiating and terminating tasks, and choosing among appropriate strategies for achieving a user's goals. CALO actively seeks out new opportunities for meeting user goals and the information it needs fully to take advantage of those opportunities. Ultimately, CALO will be trusted to act on behalf of the user in many circumstances. Interaction with the user will be primarily in terms of high-level goals, decisions, and activities.

EVALUATION, TRANSITION, AND ACQUISITION OF A SYSTEM THAT LEARNS

With traditional engineering projects evaluation can be done in a straightforward manner, determining if the documented requirements of the system have been met. Agent-based capabilities and other network centric capabilities (e.g., web services) complicate matters because the environment that they will operate under constantly changes. Add to that complication, the ability to learn new capabilities, and testing whether or not a new agent is ready to be deployed becomes a problem beyond the current body of practice.

This document lays out the problem in such a way as to identify the key issues for evaluation. By doing so, research can be targeted for the problem and solutions found. An initial experiment design is proposed as well to examine the role that evaluation will play towards transitioning cognitive systems that learn into the military environment.

CONTRIBUTION OF THIS WORK

This report describes a proposed environment called the *PAL Boot Camp*. Although this environment does not yet exist, by describing it, the primary measurements that will be necessary can be discussed in detail as these will be the key to the solution. In addition, others may be able to propose solutions to some of the problems described within.

IDENTIFICATION OF THE PROBLEMS

PROBLEM 1: WHEN IS A PAL READY TO BE FIELDDED?

A system is ready to be fielded when it has passed test and evaluation (T&E). For the military this is typically done in stages. First the system is evaluated for technical correctness and then it is introduced into an operational environment to see if the training and concept of operations allow war fighters to make good use of the system with safety. Essentially, the system is first compared to the requirements laid out in the program, then it is evaluated in an exercise setting to validate that bringing in a system that meets the documented requirements actually helps as much as anticipated.

This cannot work for a PAL for two reasons.

1. A PAL is not intended to successfully perform capabilities until it has learned the tasks involved within the operational setting.
2. The list of capabilities that a PAL will be able to do is not more than partially known until it has entered the environment and begins learning. Even then, the list is expected to grow. The capabilities theoretically could be added infinitely, though there are obvious limitations based on the resources available to the PAL.

PROBLEM 2: WHAT MUST A PAL KNOW IN ORDER TO LEARN CAPABILITIES IN THE FIELD?

If we accept that we cannot evaluate a PAL for operational capabilities prior to its use, we must still be able to determine when a PAL is ready to be sent out into the field. It is hypothesized that a PAL must have some amount of knowledge about the domain it is entering in order to learn within that domain. Two thresholds must be surpassed. First, the PAL must know enough to learn from the user and from observation. Second, the PAL must be useful enough in order for a human to be willing to have it around.

Further it is hypothesized that knowledge in other related domains will aid a PAL in learning within the domain it is to operate in. This ability, known as *transfer learning*, is important both to the introduction of a PAL into the operational environment, and in its ability to quickly respond to surprise situations and new demands by the human it is meant to assist.

PROBLEM 3: CAN A PAL GO THROUGH SYSTEMATIC TRAINING AND HOW WOULD WE MEASURE THE RESULTS?

One solution to *Problem 1* is to set up a controlled and measured training process. It is hypothesized that such a process can be set up using simulation systems that are used to train humans on their role in the operational environment. In order to accomplish this, we will have to show that the right information can be made observable to the PAL. Additionally, we will need to develop measurements to determine if a PAL has sufficient background knowledge to enter this training successfully and we will have to determine when enough training has been achieved for the PAL to graduate.

PROBLEM 4: CAN WE IDENTIFY THE CORE KNOWLEDGE NECESSARY TO A PAL?

If a PAL is to succeed in the real world, it must have the right knowledge to allow it to observe, understand, and learn from the environment. There are two basic ways we can try to identify the core knowledge necessary:

1. Analysis
2. Observation

There are advantages and disadvantages with both of these approaches. Analysis can be wrong, but observation isn't useful until the PAL gets out into the real world. It will be analysis that is necessary to give the PAL the initial knowledge it will need to learn. Observation will allow us to determine what knowledge was actually useful, if we can adequately determine it given the variability of the environments that a PAL will encounter.

THE BOOT CAMP MODEL

WHY A BOOT CAMP?

People are trained before entering new environments in the military. The basics are taught at a *boot camp*. Similarly, staff officers are trained in processes such as *Crisis Action Planning* before they join a unified command. Since crises almost by definition are not very predictable, most of what officers learn is on-the-job training. That is true of many of the knowledge-based jobs in the military, but still it is found useful to train, and in some cases, test the knowledge of individuals before they go into such environments. This ensures that the background knowledge needed is in place to allow a person to learn quickly in their new job.

A PAL faces the same challenge with the added complication, that if it is not found useful in the field, it will not be used and therefore, will certainly not learn. This chicken-and-egg problem can only be solved by ensuring that a PAL has enough knowledge to allow it to be effective enough to be utilized while it learns, so that it can improve its performance.

One solution is to immerse a PAL into a similar training environment, or perhaps the same training environment that humans are trained in. This training environment will serve the same purpose as for humans; namely, to prepare the PAL to enter the operational environment as a useful participant.

But the PAL has an advantage over humans. Once a single PAL has been trained, it can in essence be cloned. The knowledge held within one PAL can be used to initiate the knowledge in others. Therefore, once we have trained a single PAL sufficiently to operate in a crisis action planning environment, we would never have to do that again, and we can focus on a new domain. The key is measuring when a PAL has learned a sufficient amount. If we can measure PAL performance in

operational use and relate that back to the training, we can determine if there is a benefit to running further training sessions in a domain for the next generation of PAL. By observing PAL in multiple domains, we can also determine the key knowledge to enable learning over our universe of interest.

In the next section, we will describe a particular training environment for a PAL. After we do that, we will describe a global PAL training environment as it might look in an ideal world. Once this full model of PAL training is described, we will explore the details of the measurements that would be needed to make it function properly.

A PARTICULAR BOOT CAMP DESCRIPTION

The key functions identified for a PAL by SRI Inc. (leading development within the PAL program) include:

- Organize and Manage Information
 - Manage e-mail, documents, and web information
 - Organize information by tasks and user activities
- Prepare Information Products
 - Prepare meeting and event information packages
 - Organize and assemble reports and summaries
 - Draw briefing elements from email
- Observe and Mediate Interactions
 - Monitor meetings, email threads, and chat
 - Record meeting discussion, events, and action items
 - Infer tasks from email
- Monitor and Manage Tasks
 - Organize and monitor task execution
 - Monitor due dates and perform time management
- Schedule and Organize
 - Schedule meetings, events, and tasks
 - Organize task dependencies and preconditions
- Acquire, Allocate, and Optimize Use of Resources (e.g., equipment, facilities, and people)

While the natural response would be to simply test for these capabilities, the fact is that PAL will not be able to perform most of them when first provided to a user. The types of tasks to be inferred and monitored, the types of reports being generated and the sources being used for those reports, and many other details will be quite different between users. This in itself doesn't mean that the PAL is not ready for deployment. Somehow we must be certain that it can succeed at learning in the environment it is placed in.

COMMAND WORLD

During FY04 and FY05, the Space and Naval Warfare Systems Center San Diego (SSC SD), in conjunction with the Naval Postgraduate School (NPS) and SRI, conducted a series of experiments

called *Command World*. Command World was a simulation of a Crisis Action Planning (CAP) process executed by military officers playing staff officer roles.

Among the results from Command World, was the knowledge that we could use such an environment to stimulate a PAL in the areas listed above. During the CAP process, much of the information exchange was conducted using email and chat, which is easily observable through instrumentation. The primary tasks of the personnel involved were information product development providing ample opportunity for observation and assistance by a PAL.

Tens of thousands of data events were recorded in the collective Command World experiments. That data is being used by machine learning researchers to train and test their technology. However, Command World didn't provide a sufficient opportunity for a PAL to learn in order to come to conclusions regarding a particular collection of software for a version of PAL. The exercises were of short duration, and the players could change between iterations meaning that both general and personalized training of the PAL was not possible. What is needed is a structured training environment where a PAL is in use for a long period of time and where the humans using PAL provide a sample adequately covering the range of users who will be found in the environment being simulated. Specific details of the Command World series of experiments can be found in [Wong+ 2006] and Luqi [2004].

COMMAND WORLD – BOOT CAMP STYLE

If we thought of Command World as an opportunity to train a PAL, track what it had learned, and evaluate the value of its knowledge, it would be set up differently. In this section we will discuss a new Command World with those as its goals.

In our new Command World, we need day-to-day operations for a considerable time. PAL must observe and participate in a large number of basic activities in order to learn. Since PAL is not ready for operational use, we need a simulated environment, but one where the game lasts far longer than the three day Command World games played so far. Likewise, we will need some idea of what tasks we want PAL to learn about during its training, and we need to be able to instrument the system to see what knowledge it uses during reasoning about those tasks.

The first of the new Command World experiments will use the Joint Semi-Automated Forces (JSAF) simulation system using a continuation of the same scenario used in previous experiments. An interface exists to allow tactical reports to flow from JSAF to the Composeable FORCEnet (CFn) command and control capability, so users will interact with the operational data through CFn and use CFn for geographic collaboration.

The basic structure of the experiments will be as follows:

- In a pre-simulation phase, a task analysis will be done and the (estimated) minimum necessary ontology will be provided to PAL in order to bootstrap learning. (This is discussed further below in the section on *matriculation*.)
- Problems will be posed through the JSAF environment. Realistic tactical information is passed to CFn, with communications instrumented so that PAL can read message traffic.
- Players will collaborate using a combination of CFn and IRIS (a user interface to PAL that includes capabilities such as email and chat). All collaboration tools and user access to information are instrumented.

- Players must compose a force to address the tactical situation in the game through JSAF, but the force composition is decided upon during collaboration. The game then continues with the outcomes reported through the tactical communications and new problems presented.

All the time that the game is being played, PAL should be learning. Following this learning phase an exam will be administered much like is done currently to monitor the progress of the learning capability being developed by PAL [Cohen+ 2005]. This is discussed further in the section on the *graduation exam*.

After the simulation and the exam are completed, PAL will enter normal operations. In our experiments, this will be a different simulation with different human actors and with game aspects that were completely untouched in the first round. Our effort will be to simulate the transition to real world use. This is when PAL would enter operational testing in our new paradigm for evaluation. Measurements of effectiveness will be made and of most interest will be the measurements of the contribution of the boot camp, which are also discussed later in this paper. It is hoped that we can produce a boot camp that will provide sufficient training to allow a PAL to learn and contribute *on-the-job*.

THE GLOBAL PAL TRAINING ENVIRONMENT

Below is a depiction of the global PAL training process. Our experiment will mirror much of what we envision for a general purpose training and evaluation process.

First, a new capability must be developed. Engineers will write code and based on analysis provide as much of an ontology as possible. Since ontology development can literally go on forever, it is assumed that after determining the essentials, the rest will be learned. It is also the case for PAL, that the ontology being engineered is for one particular domain, and that this domain is different than the possible military transition areas. This is another reason that this effort is necessarily limited in application.

Next, a PAL or other capability will enter basic training. By developing scenarios and allowing day-to-day use in a simulated military environment, it is believed that PAL can learn the basics about the transition domain. Further training can be done in specific areas of interest (e.g., crisis action planning), where scenarios exist or can be designed. In some cases, there are training plans for staff officers and other humans that can be used to create a framework for PAL training.

After graduation, comes the “mind meld”. This phrase commonly borrowed from the television series Star Trek, indicates transfer of knowledge without the need of overt communication. In some cases, what is learned in one PAL can be easily transferred to all others, obviating the need for every PAL to go through the same training. Until the training is upgraded, or new topics are inserted, no other PAL need be brought through the same scenarios, once we are satisfied by the results. Until that point, we can continue to iterate through the training until the measurements discussed below are satisfactory. At that point the mind meld to all PAL capabilities can be made.

Finally, we enter the workforce. An instance of PAL will be paired with a person (though it is also envisioned that PAL might be fielded to support a role rather than an individual human), or more likely a group of PAL will be fielded with a group of people. Other learning systems may be fielded in different ways. Nevertheless, on-the-job training commences. PAL with, or without its particular human can go through further training, but it is envisioned that once a PAL is matched with a human, the two are inseparable.

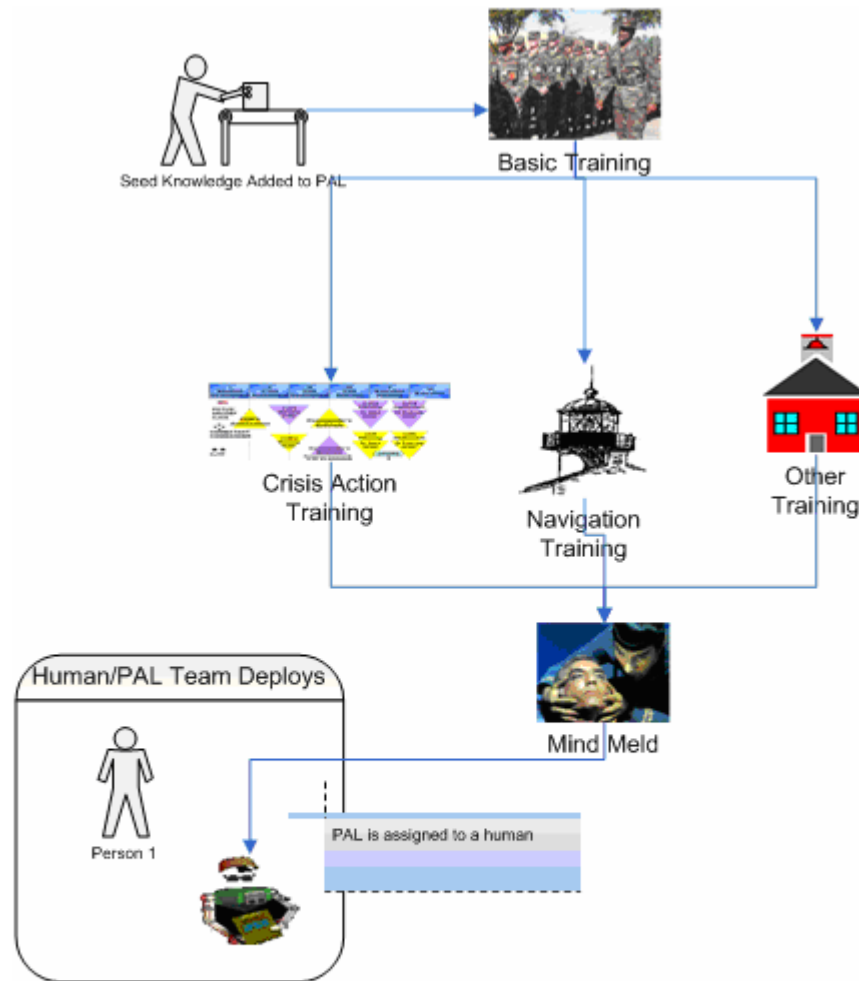


Figure 1 The Global PAL Training Process

MEASUREMENTS

There are several areas that must be measured in the processes described above. In this section, we will describe the motivation for each measurement. In some cases, a proposed metric and even a description of the instrumentation are included.

We need to work backwards to adequately understand the measurements needed. This is because there is a feedback from the operational environment that must inform the school environment. So despite the fact that the first PAL must go through the steps in order, once it has emerged from training and is in the operational environment, all future cycles begin with information generated from measurements in the field.

QUALITY OF TRAINING

The fundamental question is: when is a system that learns ready to be fielded. The descriptive answer is that it is ready to be fielded when it performs well enough to be more of a help than a hindrance and can learn quickly enough while on the job to justify the effort expended by the humans to help it learn. Therefore, the first useful measurements come from the field.

Ability to Perform

Since we are unable to work from a hard specification, we must work from a comparison of the performance of the human before and after being teamed with a PAL. Ultimately we are hoping to predict this, but early in the deployment testing of a PAL, we will need to measure human performance both with and without PAL. There are two important aspects to measure: effectiveness and efficiency. Ideally, we would like to see that effectiveness is similar (within some specified deviation from the mean performance) when a PAL is introduced as it was prior. It should be expected that efficiency will initially decrease since the human will have to spend some time informing PAL.

Idealized Effectiveness




Figure 2 Idealized Effectiveness Curves

Idealized Efficiency




Figure 3 Idealized Efficiency Curves

If this behavior is in fact realized, then specifying the operational readiness of a PAL becomes an exercise in determining the acceptable variation within effectiveness to target for initial operations, an acceptable reduction in efficiency expressed relative to prior human efficiency, and threshold values to determine if PAL performance has deteriorated or failed to improve sufficiently to warrant continued use. Operational evaluation then becomes an exercise in determining if this behavior appears within the variations specified when PAL is introduced into the operational environment.

Contribution of Knowledge

Of course, PAL could simply be improving effectiveness and efficiency because the human had to contemplate how the job is done in order to teach PAL or overcome its weaknesses. Therefore, we want to ensure that PAL is making a contribution itself, otherwise all that has been demonstrated is that people need more and a different style of training.

Although it is notoriously difficult to measure the value that particular knowledge contributes to a person or organization, there have been some recent advances on methods for doing this, such as the Intranet Efficiency and Effectiveness Model (IEEM) [Jacoby+ 2005]. There are several approaches possible for measuring the contribution that PAL makes to any process. The easier ones are less satisfactory, and as expected, those which would increase confidence are difficult to produce in the general case. From less difficult to more difficult these techniques are:

1. Survey the users to determine how much contribution was made by the PAL. Such methods are difficult to normalize (each respondent may interpret the levels differently), and people often answer surveys incorrectly or even dishonestly.
2. Perform a regression analysis of the correlation between performance and the number of changes to the knowledge of the system over time. This won't mean that the changes were the cause of the improved performance given that we cannot control for the learning being done by the human, but it will support other findings. [Cohen 1995]
3. Control for the most likely source of non-PAL improvement (learning by the human) by manipulating the only factor we can control. By removing PAL after some period and continuing to measure the performance by the human, we can use the regression results from #2 to determine if human learning or machine learning was the more closely correlated factor with changes in performance. This is an extension of the ablation testing done for CALO evaluation [Cohen+ 2005].
4. Ideally, we would like to trace what knowledge and actions were taken to perform every task. In that way we could see the role of PAL and the role of the human. This would require a detailed model of the tasks that can be performed, which isn't possible given that we want a PAL to learn new tasks that we perhaps have not yet envisioned. What we can do is compare what a human might do and what the human/PAL team did in performing a task. This is similar to what was demonstrated in [Wallace 2003]. However, this will only be feasible for a limited set of tasks, though it can help explain and validate the results from #3.

Measuring the Ability to Learn and Contribution of the Boot Camp

This is the measure of the value of the boot camp itself. Since we are not concerned with meeting hard specifications, what we want to demonstrate is that the boot camp allows PAL to learn the tasks necessary for the environment it is being placed in. We also want to determine when the learning produced in the boot camp is enough for the learning to take off in the operational context.

Measuring the ability to learn in this situation is similar to the transfer learning problem [Marx+ 2005]. Here we want the PAL to function and learn in the operational environment based on training in the boot camp. In some cases, the tasks will be the same and will not seem to be a true transfer problem, but it is expected that in the operational environment, tasks will be performed differently, and indeed this is likely to be true from one command environment to another.

To measure the ability to learn in the operational domain, and in particular the contribution that the boot camp provides to this ability, we need to measure the difference in learning between PAL that attend the boot camp and those that do not. What we want to see is if:

1. PAL that attend boot camp perform better than those that do not attend the boot camp at the beginning of their operation.
2. PAL that attend boot camp learn faster than those that do not, at least in some initial period.

Using statistical methods developed by Bamber in [Bamber 1979] and later (as yet) unpublished work, we can compare not only the current capability, but separately, the relative improvement of capability between a PAL and another PAL that has a head start due to training. Therefore, we can measure the benefits of the boot camp and also look to improve a boot camp that might be falling short. We can also get measures for rates of capability improvement in order to predict when PAL

will make a sufficient contribution to operations and directly influence decisions on suitability of the new capability.

THE GRADUATION EXAM

Graduation from a school usually means that the institution believes that a student has achieved sufficient knowledge to enter a domain where the education will be used. For PAL and other such capabilities, graduation from boot camp must mean something similar. If we have measured the results from initial attempts using the metrics above, we have some idea as to how we are doing in this one domain, from this particular boot camp. The graduation exam needs to be predictive of how a PAL or other capability will perform even when we cannot do the measurements above as they could be expensive.

For the PAL program, tests have been created using the methods employed in developing standardized exams for human students. First a task analysis was conducted, and then necessary skills and knowledge were determined. Finally, exam questions were generated to test for the necessary skills and knowledge [Cohen+ 2005]. PAL is meant to interact with human users, so using test techniques that are applied to humans seems appropriate. It might not be for other capabilities.

For the boot camp, occasionally determining the quality of graduation exams by comparing them to the results that measure the quality of training should be useful. It is not guaranteed that what is learned from the training in one domain will necessarily inform every graduation exam, but it is likely that useful lessons can be learned.

READINESS TO MATRICULATE

PAL will not even be ready to be put in the boot camp environment “right out of the box”. PAL is being developed and trained in an office automation environment as part of the research and development, but will need at least some initial knowledge about the military and command and control to perform in the boot camp. For instance, in an experiment on transfer learning (results are discussed in [Marx+ 2005] but this aspect is not), an initial ontology describing military command and staff position relationships was necessary to allow learning to occur in appointment acceptance, whereas the knowledge was already present about project management and academic staff relationships.

For the PAL boot camp, we will be producing an initial ontology, and it appears that some form of task analysis is essential to determine if the baseline ontology will be sufficient to allow learning to occur. Similar to the discussion of the graduation exam, correlating the results from this analysis in some way to graduation exam results should help inform the analysis process. Again, this is not unlike what is done with human students, as acceptance criteria is (at least advertised as) partially an attempt at predicting success towards graduation.

SUMMARY

Deciding whether or not a capability belongs in operational DOD use must change once we introduce learning into our systems. Therefore, software and indeed systems engineering practices will need to evolve. Rigid specification will have to yield to statistical methods that will indicate relative contributions and the speed at which those contributions will improve.

As we change the evaluation paradigm, we will also need to add a step to our processes that blurs the boundary between development and evaluation. Training the learning capabilities in the relevant

domains will to some degree be separate from the engineering of the capabilities and yet still is part of development. But the most valuable learning needs to take place near if not in the operational domain. If we can assume that like people, these tools can be more productive if initially trained outside the operational domain so that they are capable of on-the-job learning, we must establish this sort of capability. Thus the PAL boot camp is conceived.

REFERENCES

- [Bamber 1979] Bamber, D., “State-trace analysis: A method of testing simple theories of causation”, *Journal of Mathematical Psychology*, 19, 137-181.
- [Cohen 1995] Cohen, P., *Empirical Methods for Artificial Intelligence*, The MIT Press, Cambridge, MA, 1995.
- [Cohen+ 2005] Cohen, P., and Pool, M., “The CALO 2005 Experiment Data Analysis”, <http://calo.sri.com>.
- [Gunning 2004] Gunning, D., “Beyond Science Fiction: Building a Real Cognitive Assistant”, DARPATech 2004.
- [Jacoby+ 2005] Jacoby, G., and Luqi, “Critical Business Requirements Model and Metrics for Intranet ROI”, *Journal of Electronic Commerce Research*, Vol 6., No. 1, 2005, pp. 1-30.
- [Luqi 2004] Luqi, “Simulation Based Evaluation for Next Generation Intelligent Systems”, Technical Report NPS-SW-04-002, August 2004.
- [Marx+ 2005] Marx, Z., Rosenstein, M., Kaelbling, L., and Dietterich, T, “Transfer learning with an ensemble of background tasks “, NIPS 2005 Workshop on Inductive Transfer: 10 Years Later, 2005.
- [Wallace 2003] Wallace, S. “Validating Complex Agent Behavior”, Dissertation, University of Michigan, 2003.
- [Wong+ 2006] Wong, L. and Lange, D., “Command World”, Submitted to CCRTS 2006.