Proceedings of the 2003 Winter Simulation Conference
S. Chick, P. J. Sánchez, D. Ferrin, and D. J. Morrice, eds.

ASSESSING TECHNOLOGY EFFECTS ON HUMAN PERFORMANCE THROUGH TRADE SPACE DEVELOPMENT AND EVALUATION

Christopher R. Hale

SAIC
4031 Colonel Glenn Highway
Dayton, OH 45431 U.S.A.

ABSTRACT

Constructive simulation provides an exploratory environment for performance – effectiveness tradeoffs. However, technology trade spaces comprise many potential experiments, each containing a large sample space of experimental outcomes. Exploration of this entire space is an intractably large problem. We describe a methodology that focuses analysts only on regions of the trade space holding the most promise for effective analyses. Our methodology uses an iterative process to define the trade space, develop system and operator descriptions, parameterize the trade space and analyze performance against requirements. Each step is briefly described through the use of a notional attack aircraft crew system example. Four vectors through the trade space are identified to guide definition of specific issues modeled within the Combat Automation Requirements Testbed (CART) environment. CART constructive simulations serve a critical role by allowing rapid development and testing of alternative technologies in each area of interest.

1 INTRODUCTION

Evaluating the effects of technology in advanced systems effectiveness is a challenging problem, made more so when system effectiveness is moderated by human performance considerations. A promising way to address this challenge is through constructive simulations. These can provide exploratory environments that allow analysts to identify factors critical to performance – effectiveness tradeoffs. Merely constructing simulations, however, does not solve the technology evaluation problem. Constructive simulation, by its nature, allows great flexibility in technology combination, level of analysis and dimensions of evaluation. These factors combine to create a trade space that can be unmanageably large. One solution is to bound the trade-space in a principled manner. Doing so allows technology evaluations focused on the most relevant aspects of a technology investment program. Our method is based on the Technology Identification, Evaluation and Selection (TIES) methodology articulated by Mavris and Kirby (1999). TIES enables technology trade-offs in the early stages of aircraft design by relying on accepted engineering models of, for example, propulsion, materials and aeronautics. Reliance on such models allows an engineering design team to assess the role of technologies on predicted aircraft performance through their effects on structures, wing sizes and loadings, and propulsion as well as on well-behaved physical parameters like lift and drag. While the TIES methodology seems well-suited to technology prediction in traditional engineering domains, it exhibits shortcomings when applied to human-system integration applications.

First, few engineering models exist for human performance that lend themselves to predictions of technology effects. There are several cognitive architectures currently used to model behavior in complex operational domains (Anderson and Lebiere 1998; Rosenbloom, et al. 2000; Zacharias, et al. 2000). However, some of these architectures rely on human performance data that are somewhat controversial or require modeling at such a low level that their use for early-stage technology evaluations is limited. Others manage these problems by either "turning off" behavioral functions or limiting their levels of analysis to broad behavioral aggregates that makes informative technology evaluation difficult.

Second, critical portions of the TIES methodology rely on expert judgment. For example, development of a technology impact matrix (discussed in detail later) is based on consultation with subject matter experts in sub-disciplines of aircraft configuration as well as in the technologies under evaluation. These experts rely on both their own analysis and on disciplinary models and historical data in making predictions of performance changes correlated with each technology. Because these models are deterministic the predictions are more reliable than would be the case if the models were stochastic, as is the case in the human behavior representation community.

Furthermore, few of the parameters used in accounting for human performance are sufficiently well-behaved to support performance predictions in the presence of new technologies, further weakening TIES. We attempted to address these shortcomings by integrating the TIES methodology with CART's task network and human performance modeling capabilities. Our methodology relies on a series of principled analytical steps to define a subset of the possible trade space that represents the most informative technology evaluation possible. Each step in the methodology is designed to constrain one dimension of the trade space. Once identified, these are then combined into an evaluation environment and represented within CART. Technology alternatives are then compared against a baseline and each other to assess overall impact on crew-system effectiveness. Our methodology is summarized in Figure 1. In this section we briefly describe the methodology at a high level. Subsequent sections will discuss each step in greater detail.

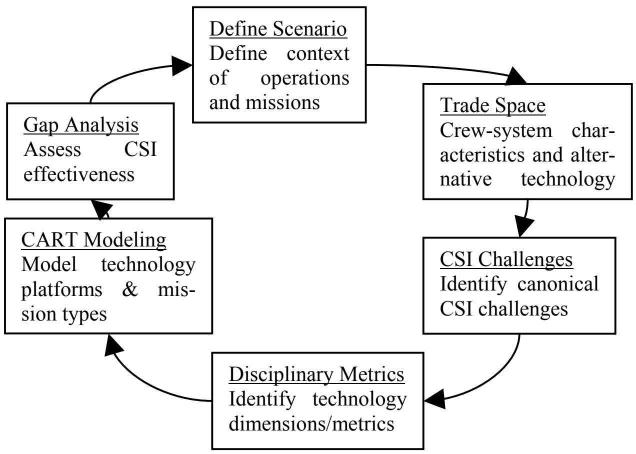

Figure 1: Methodology Overview

We begin with development of a general scenario and a set of excursions capturing plausible operational variations. The scenario provides a structure within which the technology evaluation is situated, necessary to ensure predictive validity when considering crew performance. We then develop platform alternatives. Technology evaluations of human-system interaction require developing alternative platform concepts because the constraints provided by platform configurations interact with the requirements of missions to determine human performance.

For example, crew performance in a low-altitude, subsonic aircraft versus a high-altitude supersonic aircraft is likely to vary considerably across different mission scenarios. Our next step is to define the trade space by developing a morphological matrix. This provides a structured method of identifying technology combinations from among many potentially useful candidates. The universe of possible technologies is constrained by including only those that relate to the fundamental characteristics of the system under evaluation. Since the fundamental characteristics of a propulsion system differ from those of a crew-system technology, alternatives appearing in respective morphological matrices also will differ. When candidate technology alternatives have been identified, they are aggregated into a set of platform alternatives. These are combined with mission excursions into a CSI challenge matrix. Challenges specific to platform-excursion combinations are then identified and summarized across the challenge space using a set of rules to be discussed below. Disciplinary metrics are then identified for a set of canonical CSI challenges. This information is combined with information in the morphological matrix to guide CART model development. The final step is to conduct gap analyses comparing each platform to baseline technology and to each other. The remainder of the paper discusses each step of the methodology.

2 STEPS OF THE METHODOLOGY

2.1 Scenario/Excursion Construction

Our method begins with construction of a base scenario and excursions from this base. Scenario construction serves as the fundamental problem definition activity within which technology evaluation is situated. This establishes the context and boundaries of subsequent technology evaluations. Our scenario construction follows three standard steps. First, we envisioned a standard geo-political background leading to a plausible military confrontation. Second, the background context was populated with technologies assumed to represent operational challenges for the system under study. This can range from current technologies to those expected to be operational at some point in the future. Third, several mission excursions were defined to help identify a range of requirements for the platforms under study. Platforms might, for example, be required to carry out missions ranging across threat suppression, time-critical targeting (TCT) and attack of hardened tunnels.

2.2 Morphological Matrix for Technology Concept Identification

The scenario for our example consists of airborne attack operations. Other scenarios might include commercial cargo delivery, ground assault command and control or counter-terrorism intelligence analysis. Whatever the specific scenario, the next step defines the trade space of evaluation. Such a space is formed by combining characteristics of the system of interest with technology alternatives that potentially address the characteristics. Obviously, these characteristics will differ with both the scenario and the system of interest. Table 1 shows an example of how one might construct a morphological matrix for an airborne attack crew-system. The left side of the table identifies 6 characteristics: flight control, command and

Table 1: Morphological Matrix for a Conceptual Attack Aircraft

Crew system characteristics	Alternative 1	Alternative 2	Alternative 3	Alternative 4	Alternative 5
Flight control	Fully automatic	Auto-planned; manual un-planned	Auto cruise, targeting, and known threats; manual response to pop-up threats	Auto cruise; manual other	Fully manual
C4ISR	Fully pre-filtered from MC2A	Targeting from MC2A; threat data directly from ISR assets	Targeting and known threats from MC2A, pop-up threats directly from ISR assets, auto-BDA	Targeting and known threats from MC2A, pop-up threats directly sensed, manual BDA	Raw data directly from ISR assets
Threat management	Fly faster than threats	Highly lethal Self-defense	Reliable early detection and avoidance	Stand-off	N/A
Target lethality	Weapon volume	High weapon accuracy	Weapon variety	Coordinated weapon delivery	N/A
Detectability	Stealthy, high altitude, supersonic	Stealthy, low altitude, supersonic	Stealthy, high altitude, subsonic	Not stealthy, high altitude, supersonic	Not stealthy, low altitude, supersonic

control, situation awareness, location/orientation, safety and lethality. We then combine these with means of realizing each characteristic. For example, alternative means of achieving flight control include fully automatic control, fully manual control and methods utilizing these extremes in different combinations. Selecting one alternative from each row of the trade space defines composite technology concepts for further exploration. Although it is possible to generate a large number of concepts from a morphological matrix the number of plausible concepts typically will be smaller than the factorial combination of all tabled attributes. Furthermore, constraints imposed by the scenario and the state of technology for the time frame envisioned will eliminate some alternatives from consideration. For example, flying faster than threats is not likely to be a plausible means of realizing threat management, thereby making this alternative an unlikely participant in further evaluations. The value of the morphological matrix is in characterizing the entire trade space from which specific technological alternatives can be selected for evaluation.

2.3 Identify Crew-System Integration (CSI) Challenges

Construction of a morphological matrix allows analysts to define a relevant trade space and concepts within that space for evaluation. We next define CSI challenges for the selected concepts to be evaluated against. We formulate the table by combining mission requirements from the scenario with technology concepts from the morphological matrix. Table 2 provides a notional example. The content of this table typically is developed by means of an analytical process involving experts in mission requirements, technology characteristics and crew-system integration. The process is guided by consideration of characteristics identified when building the morphological matrix. For example, concepts aggregated into a standoff attack platform exhibit specific challenges in the areas of situational awareness, safety against threats and target lethality when executing threat suppression missions. As these challenges are identified they are placed into the appropriate cells of the challenge matrix. It should be noted at this point that not all cells in a CSI challenge matrix will necessarily be completed. It is possible, in fact probable, that some platform concepts will be inappropriate for particular missions. For example, standoff platforms are not likely to be appropriate for time-critical targeting.

Note that Table 2 also contains marginal cells representing summary information across mission types and platforms. Summarizing across missions and platforms allows analysts to identify subsets of CSI challenges affecting multiple platforms and missions. Further summarizing across the marginals produces a canonical set of challenges that guide subsequent task network modeling. Creation of the summary platform/mission challenges or the canonical challenge set is not simply a matter of enumerating the contents of cells in the table. Rather, we have discovered several rules that are useful in guiding the process.

First, challenges that appear consistently across cells should be captured. Second, technology availability should be considered. CSI challenges that seem satisfied by currently available (old) technology should be eliminated from further consideration. Third, mission phase criticality affects the "weighting" of some CSI challenges. Since some mission phases are arguably more important than others (e.g., threat management over cruise) the challenges associated with the former phases should be preferred in the summarization process. Fourth, challenges that are related to brittle technologies should be retained An example of a brittle technology would be that of communication. Although current communication capabilities are impressive, they also can be easily disrupted. Disruption would lead to potential mission degradations that should be explored with new technology concepts. Fifth, the perceived cost of new technologies should be considered. While it is possible, for example, to imagine particle

Table 2: CSI Challenges

Mission Requirement	Stand-off	Subsonic	Supersonic	**Mission Challenge Summary**
SEAD	• Weapon guidance over the horizon must be highly accurate • Multiple targets with coordinated attack management	• Protect against passive detection • Merge SAR, EO, tactical and human visual information	• Target coordinate updates will be crucial • Target imagery	• Target location updates • Multiple image sources • Weapon guidance and control
Time-critical targeting	N/A	• Maintain SA in real-time environment • C4ISR integration	• Target updates must be timely	• ISR integration
Hardened targets	Weapon guidance over the horizon must be highly accurate	• Must have accurate target coordinates at point of release	• Weapon guidance on-board?	• Target coordinate accuracies • BDA
Platform Challenge Summary	• Over the horizon BDA • Long range weapon guidance • Target/coordinate updates	• Pop-up threats • Imaging effectiveness • ISR integration	• High speed weapon delivery • Mission update timeliness	Dynamic re-planning Image fusion Synthetic/enhanced vision Human-system integration

beam weapons as potential solutions to some challenges the cost of such technologies is unlikely to make them viable within a reasonable time frame. Sixth, challenges that might be associated with technologies having particularly high payoff should be explored. For an attack scenario like that used in this example, one might identify the areas shown in the lower right cell of Table 2 as canonical challenges. These then serve as the primary focal areas carried forward into the simulation phase.

2.4 System Effectiveness Metric Identification

Development of the morphological matrix enables identifying aggregate technology concepts. The CSI challenge matrix identifies areas critical to mission success. Our next step is to relate this information to human performance by constructing a technology impact matrix (TIM), an example of which is shown in Table 3. This follows a 3-step process: (1) technology dimensions of each canonical CSI challenge are identified, (2) metrics are identified for each technology dimension and (3) factors affecting the metrics are formulated. Referring to Table 3, one dimension of dynamic mission planning is the method of uplinking and downlinking information. Three metrics related to uplink/downlink method are time needed to input new information into on-board automated systems (update time), input accuracy during mission updating, and the ratio of manual to automated updating required (update efficiency). Disciplinary factors affecting these metrics include crew workload, goal priorities, stress and message perceivability. A separate TIM is defined for each dimension of each canonical challenge. Note that technology dimensions, combined with platform concepts, constrains the technologies available for consideration during this phase of evaluation. Maintaining this constraint is important in preventing a run-away consideration of all possible technologies, though one can always go back to the morphological matrix to formulate new concepts if desired.

2.5 CART Modeling

The CART environment provides a way to empirically conduct the evaluations needed to make technology investment recommendations. CART is a goal-oriented, task network modeling tool based on the IMPRINT discrete event simulator. It allows analysts to describe the structure of mission execution to an arbitrary level of detail. Beginning with goals, the structure of a task environment is decomposed as a set of functions and tasks through as many levels as needed to accurately capture the human-system aspects of task execution, as shown in Figure 2. The goal states provide the model organization and control that represents the adaptive nature of human behavior. When representative goal states and basic behavioral task structure have been defined, tasks are populated with information concerning task times, standard deviations and accuracy, as exemplified in Figure 3.

Distributional assumptions also are included at this step. A further specification at this time that is unique to the CART modeling environment involves defining micro-models that provide internal references to human performance data. This step is where much of the specification of disciplinary parameters takes place. For example,

Table 3: Technology Impact Matrix Example

Uplink / Downlink Method		**Platform Technologies**		
Metrics	**Disciplinary Parameter Vector**	Stand-off	Sub-sonic	Super-sonic
Update time	Situation awareness	+ 3%	+ 9%	+ 4%
	Workload	+ 10%	- 2%	- 12%
	Memory load	0	+ 5%	- 6%
	Goal priority	+ 30%	+10%	- 10%
Input accuracy	Stress	+ 22%	- 15%	+ 18%
	G-loading	+ 23%	- 15%	+ 20%
	Message perceivability	+ 14%	- 20%	+ 18%
Update efficiency	Perceived goal cost	+ 13%	- 29%	+ 18%
	Message comprehensibility	+ 15%	+15%	+15%

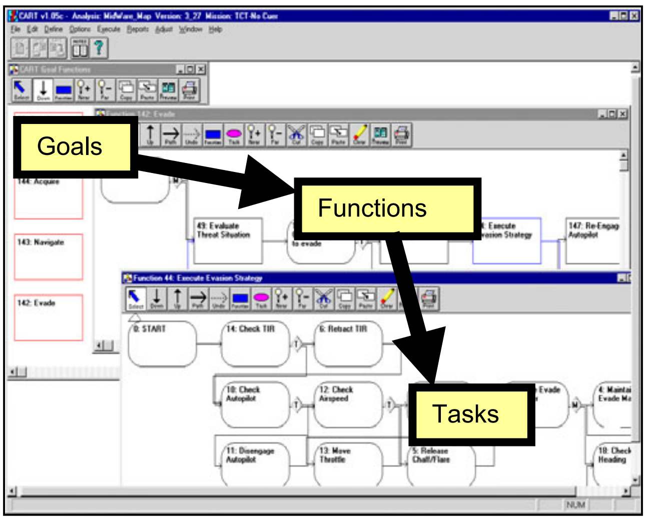

Figure 2: Hierarchical Structure of CART

specification of g-force effects on data-pad input accuracy would be contained in CART as a micro-model. The final step in basic CART model building is to specify decision nodes. As shown in Figure 4, CART allows three types of decisions. Tactical decisions test for conditions in the task environment and direct flow of control to particular tasks mapped to conditional outcomes. Multiple decisions direct flow of control to two or more tasks based on completion of a prior task. All subsequent tasks begin performing at the same time. Probabilistic decisions direct flow of control to two or more subsequent tasks based solely on *a priori* probabilities associated with a prior task.

To illustrate our approach to technology effects evaluation, assume a CART simulation consisting of two strike aircraft, each having different goals and each encountering different critical simulation events arising out

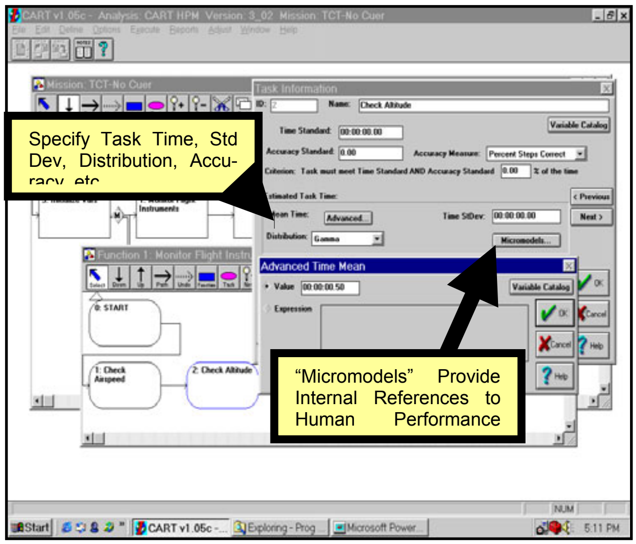

Figure 3: Task Performance Specification

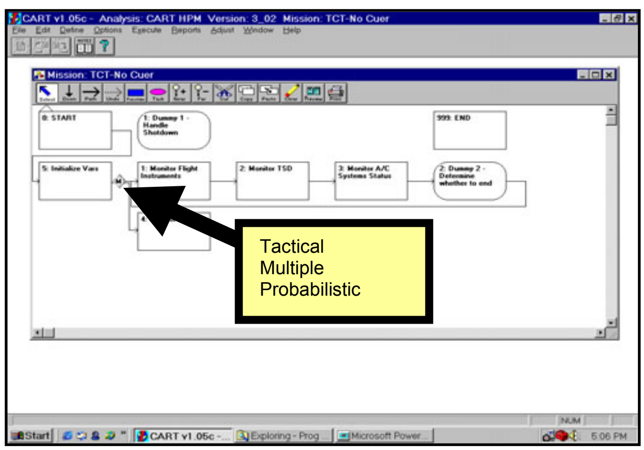

Figure 4: CART’s Three Decision Types

of interactions with other simulation entities. These three categories; focal objects, participating objects & technologies, and events; are varied to produce the conditions for evaluation. A notional simulation timeline capturing these categories is shown in Figure 5. We develop the goals for our model through analysis of historical data concerning strike aircraft operations (Schumann, et al. 1985) and based on our own experience in this domain. Examples of goals in this domain typically include maintaining geo-positional orientation, ensuring ownship safety, vehicle control and destroying assigned targets. When goals have been defined we then develop function – task descriptions (decompositions) for each critical event.

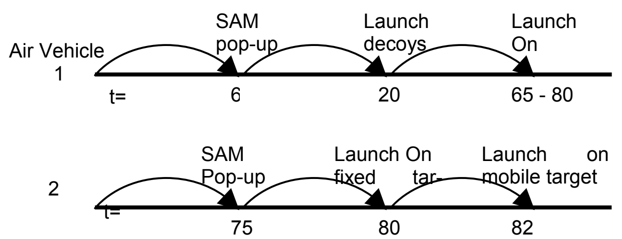

Figure 5: Notional Simulation Timeline

Figure 6 provides an example of our decomposition approach. This figure outlines a task network constructed by a notional strike aircraft in the presence of a pop-up surface-to-air missile (SAM) threat. Figure 6a contains the hierarchical task structure. Figures 6b and 6c contain

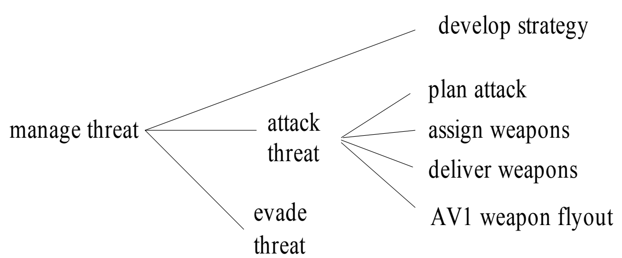

Figure 6a: Pop-Up Threat Task Structure

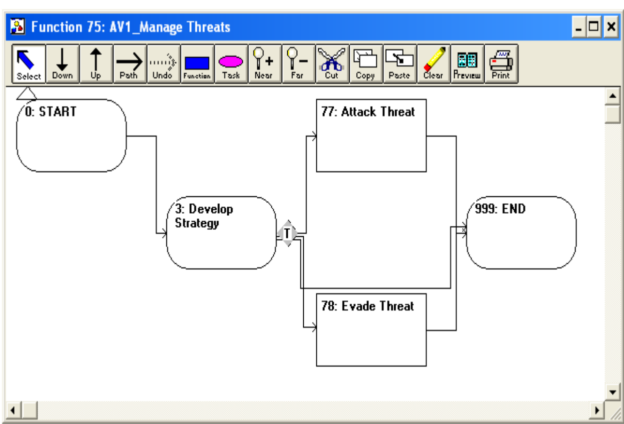

Figure 6b: A CART 'Manage Threat' Function

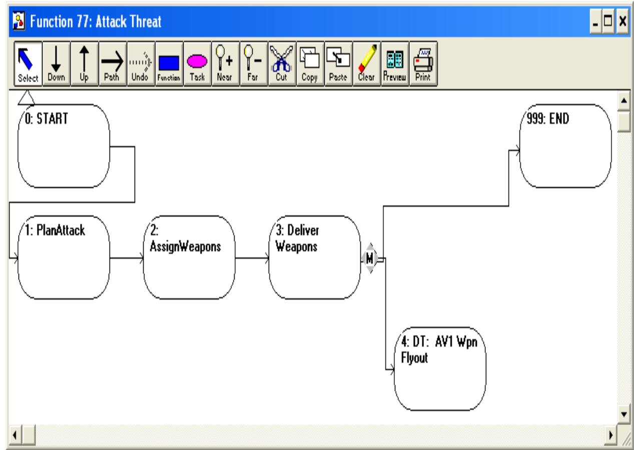

Figure 6c: A CART 'Attack Threat' Task Network

CART representations of the resulting functional decompositions for the top-level function 'manage threat' and for the second-level function of attacking the threat.

After developing the basic task networks we specify performance dynamics for each task by defining mean task times and standard deviations, task time distributions, accuracy standards and measures, triggering conditions for each task, and beginning and ending conditions. Task times are based on historical and empirical data and on estimates of subject-matter experts (SMEs).

Triggering, beginning and ending conditions are based on a rational analysis of the task requirements and are reviewed by SMEs. Where appropriate, micro-models are specified for tasks based on available empirical data and estimates of technology effects from team members and SMEs.

We then specify decisions for each of the networks. As a general rule, in the case of simulations constructed to study system performance in reactive, intelligent, often adversarial environments we have found that the majority of decisions are tactical. The decisions typically take the form of production rules, with condition clauses testing either for the presence of states onboard the strike aircraft or for changes in known states of potential adversaries and action clauses either rearranging the priorities of goals or changing the values of variables needed by the task network. All tactical decisions are included at the lowest levels of the networks.

Upon completing the basic structure of the system we define dependent measures based on the technology impact matrices developed earlier. In the current simulation, for example, two measures might be defined. Measure 1, applying exclusively to wave 1, might consist of the status of a particular SAM (viable or destroyed) at the point of overflight of the SAMs position by the wave 1 aircraft. Measure 2, exclusive to wave 2 aircraft, might be the time needed to update targeting information. This measure is taken directly from the TIM example given in Table 3.

As can be seen from an inspection of Table 3, selection of a particular metric informs the analyst of the disciplinary parameters that should be included in the CART model. Disciplinary parameters determining these metrics are built into the CART model through its network structures, goal priorities and human performance components. These include both the micro-models intrinsic to CART and external model calls to general human performance architectures such as ACT-R or specialized subroutines. For this particular metric (update time) the model should include situational awareness, workload, memory load and goal priority parameters. The first three of these are included as micro-models within appropriate tasks whereas goal priority is included as a structural aspect of the model itself. We typically construct a "stealth viewer" to collect the data of interest as each experiment proceeds. This viewer consists of a simple spreadsheet containing air vehicle data, range to points of interest, key events indexed by their time of occurrence and data regarding SAM mode and status.

The dynamics of the aircraft are altered by changing simulation parameters such as airspeed, communication lags between aircraft and controlling entities (such as an air operations center), altitude and weapon characteristics. This allows us to develop functional definitions of different platforms.

This approach allows us to describe technology effects on operator performance at whatever level of analysis is desired. In our modeling technology effects are apparent in two ways. First, by affecting task network structure a particular technology affects mission outcomes by adding, deleting or changing the ordering of network elements. For example, a mission updating technology (addressing the dynamic mission planning CSI challenge) that automatically inputs target coordinates into a flight computer might result in the elimination of several tasks on the part of an aircrew during the weapon delivery portion of a simulation event. This effect would be realized in CART as an increase in target coordinate recording accuracy, a decrease in time required to register new coordinates with the flight computer, an increase in the number of targets that can be entered into the flight computer per unit time and so on.

Second, technologies can affect the metrics that index task execution. A technology requiring manual reprogramming of target coordinates might, for example, affect coordinate update time either by increasing workload or by promoting coordinate input above target designation (thereby altering goal priority) at a critical point in the task flow. Either of these can negatively affect coordinate update times. Technologies also can affect task reliability.

Consider, for example, a technology that relies on satellites for communicating mission updates. When communication mediated by these satellites is disrupted, other methods are used. However, task execution performance can became much more variable, thereby resulting in degraded mission performance. CART allows an analyst to vary these two parameters by (1) defining mean accuracy associated with a particular task and (2) specifying the distribution and variance associated with a task.

3 COMPARATIVE GAP ANALYSIS

Our final step is to assess the mission effectiveness associated with our three evaluation platforms. To do this we model each system metric based on the vector elements of Table 4. In most cases we anticipate these models will take the following form:

$$R_p = b_0 + \sum b_i k_i + \sum b_{ii} k_i^2 + \sum\sum b_{ij} k_i k_j$$

In this form R_p represents a performance metric, b_i represents linear regression coefficients, b_{ii} represents quadratic coefficients, b_{ij} represent cross-product coefficients and each k term represents disciplinary factors from the parameter vector in Table 4. We then conduct simple comparisons across each of the platforms on the metrics of interest.

4 DISCUSSION

The methodology discussed above provides a principled way of moving from an open-ended space of potential technology concepts to predictions of technology effects on human-system effectiveness. The value of this method resides in providing: (1) a way to define a bounded trade-space within which technology alternatives can be identified, (2) a method of combining mission requirements with high-level trade-space technology alternatives to assist in identifying areas to help focus technology predictions, (3) a modeling tool to represent the task networks required to carry out mission requirements using identified technology concepts and (4) an analysis method relating technologies directly to crew-system effectiveness metrics.

Future work will be concentrated in several areas. The first addresses how to represent the technology under review. If representing technology can properly be considered a hierarchy, then the question becomes what level of analysis is meaningful. Our technology representations to date have primarily been functional and have occupied a fairly high level of aggregation, the platforms discussed in above. However, it is possible to represent technology on at least two lower levels: Technology class level and a specific system level. These levels, when applied to radar for example, might include sensing, SAR and a specific system or application. Choosing one representational level or another when executing the methodology outlined here would lead to different simulation outcomes.

A similar question arises in consideration of human performance. Again, our simulations to date have been limited to relatively high levels of analysis in our human performance modeling. While maintaining this level has been useful in comparing platform configurations and in identifying broad crew-system effectiveness issues, it is less useful in evaluating specific alternatives (e.g., SAR versus EO as a target recognition aid). On the other hand, moving the human performance level of analysis to a low level risks (1) creating a problem of proliferating process boxes and (2) disconnecting human performance evaluations from an overall understanding of system effectiveness. Worse, still, the levels of analysis problems for technology and human performance are interrelated.

A third challenge for future work with this methodology is how to account for variation in crew performance, at both the task and method levels. Crews often achieve goals by combining tasks in different sequences or by using different tasks altogether. Additionally, tasks can be accomplished through variations in methods. This variability creates problems for technology evaluation, as it is difficult to state with certainty that one combination of technology is clearly superior to another without regard to task and method variation.

Finally, the problem of technology interactions must be addressed. As has been pointed out by Overdorf (2002), technologies might lead to performance improvements when considered in isolation but to performance degradations when combined. This raises the problem of combinatorial explosion in which all possible combinations of all technology candidates must be considered. We finessed this problem in the current study by combining technologies into functionally defined platforms defined to address only the mission excursions of our scenario. However, this strategy is more ad hoc than we would like.

We feel that the methodology discussed here, particularly the addition of the CART modeling and simulation environment, holds great promise for evaluating technology effects on crew-system performance. As the challenges outlined above are addressed, the general method should greatly facilitate technology investment decisions and the place of crew-systems in simulation-based acquisition.

REFERENCES

Anderson, J. R. and C. Lebiere. 1998. *Atomic Components of Thought*. Hillsdale, NJ: Lawrence Erlbaum.

Mavris, D. N. and M. R. Kirby. 1999. Technology identification, evaluation and selection for commercial transport aircraft. Presented to the 58th Annual Conference of the Society of Allied Weight Engineers, San Jose, CA.

Overdorf, R. L. 2002. Personal communication.

Rosenbloom, P. S., W. L. Johnson, R. M. Jones, F. Koss, J. E. Laird, J. F. Lehman, R. Rubinoff, K. B. Schwamb, and M. Tambe. 1994. Intelligent automated agents for tactical air simulation: A progress report. *Proceedings of the Fourth Conference on Computer Generated Forces and Behavioral Representation*, 69–78. Orlando, FL.

Schumann, S. A., C. A. Smith, B. K. Howard, P. J. Gros. 1985. Advanced strategic aircrew systems study (ASASS): Pilot factors research for crew systems development. Technical Report No. AFWAL-TR-85-3032, Air Force Wright Aeronautical Laboratories.

Zacharias, G. L., A. X. Miao, C. Illgen, and J. M. Yara. 1996. SAMPLE: Situation Awareness Model for Pilot-in-the-loop Evaluation. Technical Report No. R95192, Charles River Analytics, Cambridge, MA.

AUTHOR BIOGRAPHY

CHRISTOPHER R. HALE is a Senior Scientist with the Aeronautical Systems Division of SAIC, located in Dayton, Ohio. His research interests include human performance modeling and the roles of technology on cognitive performance. He has conducted research in human memory, human – system integration, mental models and diagnostic reasoning. He holds a PhD in cognitive science from the Georgia Institute of Technology.