

A University Research Group Experience with Deploying an Artificial Intelligence Application in the Center of Gravity Analysis Domain

Gheorghe Tecuci, Mihai Boicu, Dorin Marcu, Bogdan Stanescu, Cristina Boicu, Marcel Barbulescu

MSN 4A5, Learning Agents Center, George Mason University, Fairfax, VA 22030, USA
{tecuci, mboicu, dmarcu, bstanesc, ccascava}@gmu.edu, mbarb@cs.gmu.edu, <http://lalab.gmu.edu>
tel: 703 993 1722; fax: 703 993 1710

Abstract

This paper presents the experience of a university research group that has successfully deployed an application of its artificial intelligence research. It identifies some of the factors that have contributed to this success, and proposes a framework for future deployment activities that are consistent with the mission of a research university.

Introduction

This paper presents the successful experience of deploying an application of artificial intelligence by a university research group. This work was part of a multi-objective collaboration between the Learning Agents Center of George Mason University, on one side, and the Center for Strategic Leadership and the Department of Military Strategy, Planning, and Operations of the US Army War College, on the other side. A distinguishing feature of this collaboration is the synergistic integration of artificial intelligence research, with military strategy research, and with the deployment of agents in education.

The artificial intelligence research objective of this effort was the development of a learning-based approach to building knowledge-based agents. The military strategy research objective was the formalization of the center of gravity (COG) analysis process. Finally, the third objective of this effort was to enhance the educational process of senior military officers through the use of intelligent agent technology. Each of these three objectives is recognized as important and difficult in its own right. Our experience is that addressing them together, in a synergistic manner, has resulted in faster progress in each of them. In particular, we have validated our research. We have also succeeded to develop and deploy intelligent agents for strategic center of gravity analysis. These agents are used in several courses at the US Army War College, since Winter 2001.

We consider that this experience offers a new perspective on how to integrate research in artificial intelligence, with research in a specialized domain, and

with the development and deployment of prototype systems in education and practice, by a university research group.

The next section provides more details on the artificial intelligence technology that was at the basis of the deployed application. The application is described in the subsequent section. Then the paper identifies some of the most important factors that have contributed to the success of this deployment, and some of the lessons learned. Finally, the last section presents a few conclusions and our plans of expanding on this success.

Artificial Intelligence Technology: Disciple Learning Agent Shell

We are researching a theory and associated methodologies and tools for the development of knowledge-based agents. The basic idea of our approach, called Disciple, is to develop a learning agent shell that can be taught directly by a subject matter expert to become a knowledge-based assistant. The subject matter expert interacts directly with a Disciple agent, to teach it to solve problems, in a way that is similar to how the expert would teach a human apprentice, by giving the agent examples and explanations, as well as by supervising and correcting its behavior. The agent learns from the expert by generalizing the examples and the explanations to build its knowledge base, and to become a better assistant (Tecuci 1998, Boicu et al. 2001; Tecuci and Boicu 2002).

The Disciple-RKF learning agent shell is the implementation of the most recent version of the Disciple approach. It includes a general problem solving component which is based on the task reduction paradigm. In this paradigm a complex problem solving task is successively reduced to simpler tasks, the solutions of the simplest tasks are found and these solutions are successively combined into the solution of the initial task. The knowledge base of Disciple-RKF consists of an object ontology that describes the entities from an application domain, and a set of task reduction and solution composition rules expressed with these objects. The learning component of Disciple-RKF integrates several learning strategies, such as learning from examples, learning from explanations, analogical learning,

and apprenticeship learning. This multistrategy learning component allows a Disciple agent to develop its knowledge base through a mixed-initiative process which exploits the complementarity between human and automated reasoning. It creates a synergism between the subject matter expert who has the knowledge to be formalized and the agent that knows how to formalize it.

The Disciple-RKF learning agent shell was used to build the deployed Disciple-COG agent, as briefly described in the following, and in (Tecuci et al., 2002a).

First, we have worked with Jerome Comello, subject matter expert from the US Army War College, to develop a task reduction-based model of how he performs center of gravity (COG) analysis. We have analyzed several representative scenarios, and have formulated the task reduction steps leading to the identification and testing of the center of gravity candidates for those scenarios. The resulting reasoning trees revealed some of the object concepts that needed to be present in Disciple’s ontology so that it can perform this type of reasoning. Using the ontology building tools of Disciple-RKF (Stanescu et al., 2003), we have developed this object ontology, a fragment of which is presented in Figure 1.

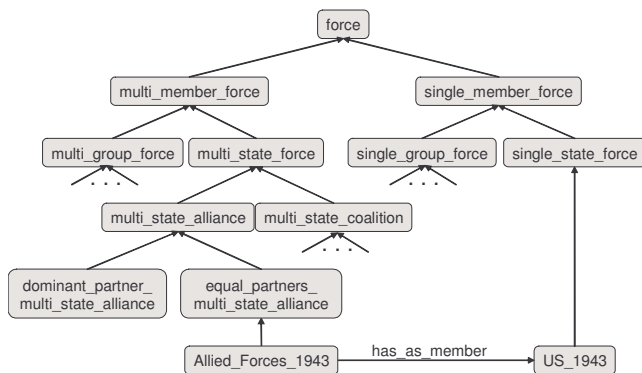


Figure 1: Fragment of the ontology of Disciple-COG.

After the object ontology has been developed, we have taught Disciple how to analyze the sample scenarios, following the reasoning trees defined with the subject matter expert (Tecuci et al., 2002b). Figure 2, for instance, shows a fragment of the reasoning tree corresponding to the World War II – Sicily scenario.

From each task reduction step (represented by a task, a question, an answer and one or several subtasks) Disciple learned a general task reduction rule, using the object ontology as a generalization hierarchy. For example rule R4, learned from the task reduction step at bottom of Figure 2, is represented in Figure 3. The top part of Figure 3 shows the informal structure of the rule which preserves the natural language of the expert and is used in agent-user communication. The bottom part of Figure 3 shows the formal structure of the rule which is used in the actual reasoning of the agent. Notice that this is a partially learned IF-THEN rule with two conditions (the plausible upper

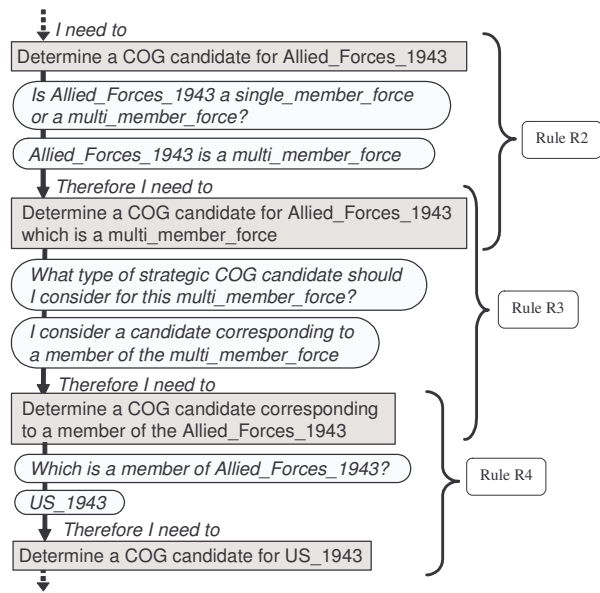


Figure 2: Fragment of a reasoning tree.

bound condition, and the plausible lower bound condition) that define a plausible version space for the exact condition to be learned (Tecuci, 1998).

As Disciple learned new rules, the interaction with it evolved from a teacher-student interaction, toward a collaboration in COG analysis. During this process, Disciple learned not only from our contributions, but also from its own successful or unsuccessful problem solving attempts, which led to the refinement of the learned rules.

The next section describes the developed Disciple-COG agent from the user’s perspective.

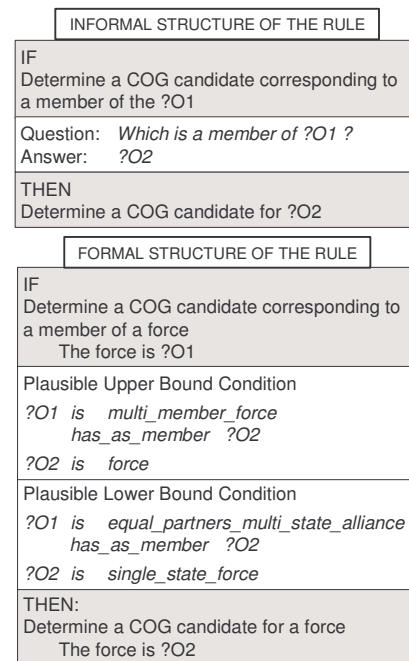


Figure 3: A partially learned task reduction rule.

Disciple-COG: A Deployed Intelligent Assistant for Center of Gravity Analysis

The concept of center of gravity is fundamental to military strategy, denoting the primary source of moral or physical strength, power or resistance of a force (Strange 1986). The most important objectives of a force are to protect its center of gravity and to attack the one of the enemy.

Center of gravity determination requires a wide range of background knowledge, not only from the military domain, but also from the political, psychosocial, economic, geographic, demographic, historic, international, and other domains (Giles and Galvin 1996). In addition, the situation, the adversaries involved, their goals, and their capabilities can vary in important ways from one scenario to another. Therefore, when performing center of gravity analysis, experts rely on their own professional experience and intuitions, without following a rigorous approach.

Inspired by the work of Joe Strange (1996), and working with experts from the US Army War College, particularly Jerome Comello, we have developed a computational “generate and test” approach, to center of gravity analysis. Center of gravity candidates from different elements of power of a force (such as government, military, people, economy) are identified during the generation phase. For instance, a strong leader is a center of gravity candidate with respect to the government of that force. Then, during the testing phase, each candidate is analyzed to determine whether it has all the critical capabilities that are necessary to be the center of gravity of a force. For example, a leader needs to be protected, stay informed, communicate (with the government, the military, and the people), be influential (with the government, the military, and the people), be a driving force, have support (of the government, the military, and the people), and be irreplaceable. For each capability, one needs to identify the essential conditions, resources and means that are required to be fully operative, and which of these, if any, represent critical vulnerabilities (i.e. are deficient or vulnerable to neutralization, interdiction or attack in a manner achieving decisive results). The candidates that lack any of the required capabilities are eliminated, the

center of gravity being among those that are not eliminated.

Since Fall 2000 we have developed several Disciple-COG agents for strategic center of gravity analysis, based on successive versions of the Disciple-RKF learning agent shell, as described in the previous section. The Disciple-COG agents have been used in all the eight sessions of the “Case Studies in Center of Gravity Analysis” course, taught at the US Army War College since Winter 2001. These sessions have been attended by 71 students (senior US or international fellows from all the military branches).

Each student works with a personal copy of a Disciple-COG agent. The agent guides the student to determine the strategic center of gravity of a force in a war scenario, helping him/her to learn a structured methodology to solve this problem. Examples of war scenarios are War World II in 1943, or the current war on terror against Al Qaeda.

First, Disciple-COG guides the student to identify and describe the aspects of a war scenario that are relevant for COG analysis. The student-agent interaction is very easy and natural for the student, taking place as illustrated in Figure 4. The left part of the window is a table of contents, whose elements indicate various aspects of the scenario. When the student clicks on one aspect, the right hand side of the window displays specific questions intended to acquire from the student a description of that aspect, or to update a previously specified description. Student’s answers lead to the generation of new items in the left hand side of the window, and trigger new questions from the agent.

The student is not required to answer all the questions

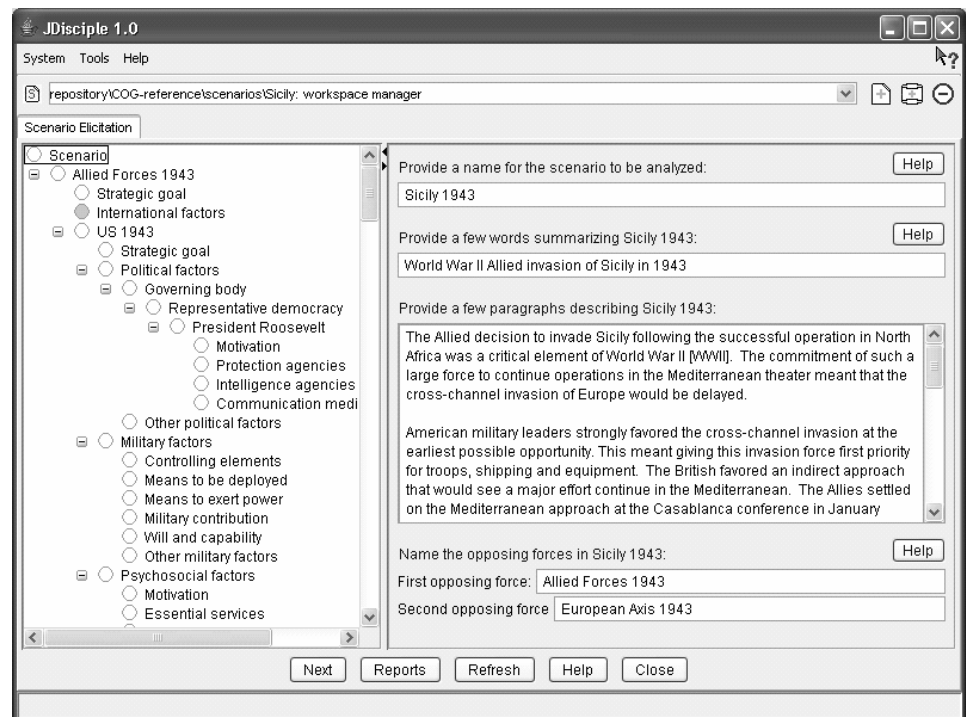


Figure 4: Scenario Elicitation interface.

and Disciple-COG can be asked, at any time, to identify and test the strategic center of gravity candidates for the current specification of the scenario. Figure 5 shows the interface of the solution viewer.

In the left hand side Disciple-COG lists the strategic center of gravity candidates for each opposing force. When the student clicks on one of them, the justification of why it was identified as a candidate, or the justification of the testing result, is displayed in the right hand side of the viewer (depending of which tab is selected by the student).

At the end of the analysis, Disciple-COG generates a report containing both the description of the scenario, and the analysis of the identified center of gravity candidates. The student then uses a word processor to finalize the report generated by Disciple-COG. He or she is required to critically analyze Disciple-COG's logic, correct or complete it, or even reject it and provide an alternative line of reasoning.

This is the first time that an intelligent agent for the strategic COG analysis has been developed.

Disciple was used at the US Army War College, while still under research and development. The "Case Studies in Center of Gravity Analysis" course (the COG course) is offered twice a year, in Winter (Term II) and in Spring (Term III). Many of the students that have taken the COG course in Winter have also enrolled in the "Military Applications of Artificial Intelligence" course (the MAAI course) in the Spring term. In the MAAI course the students act as subject matter experts, teaching personal Disciple agents their own reasoning in center of gravity analysis. This allowed us to perform

unique knowledge acquisition experiments that would have been very costly if we were to hire the subject matter experts. More details on the use of Disciple in the MAAI course are provided in (Tecuci et al. 2002a, 2002b).

We have treated each session of the COG or MAAI course as an experimentation with a new version of the Disciple-COG agent (in the case of the COG course), or with a new version of the Disciple-RKF agent shell (in the case of the MAAI course). Based on the lessons learned from each course, we have developed an expanded and improved version that was used in the next session of the course.

Factors Critical to Success and Lessons Learned

There are many factors that have contributed to the successful development and deployment of Disciple-COG. Moreover, this experience has taught us several lessons, which we consider important to mention. While the following are the most important factors and lessons learned, this is not an ordered list.

Problem of great importance to the customer:

Correctly identifying the centers of gravity of the opposing forces is of highest importance in any conflict. Therefore, in the education of strategic leaders at all the US senior military service colleges, there is a great emphasis on the center of gravity analysis. The Center for Strategic Leadership and the Department of Military Strategy, Planning, and Operations of the US Army War College have seen a great value in the development of an intelligent agent that would enhance the education of the future strategic leaders of the military, and have therefore strongly supported this effort.

Synergistic research environment:

By tightly integrating research in artificial intelligence with research in military strategy, we have involved the subject matter experts in our own research and we have involved ourselves in the research of the subject matter experts, all of us having an important commitment to this work.

Working closely with the customer: From the very beginning we have collaborated with the professors from the US Army War College, to develop a system that best suited their needs and expectations. Jerome Comello, the instructor of the "Case Studies in Center of

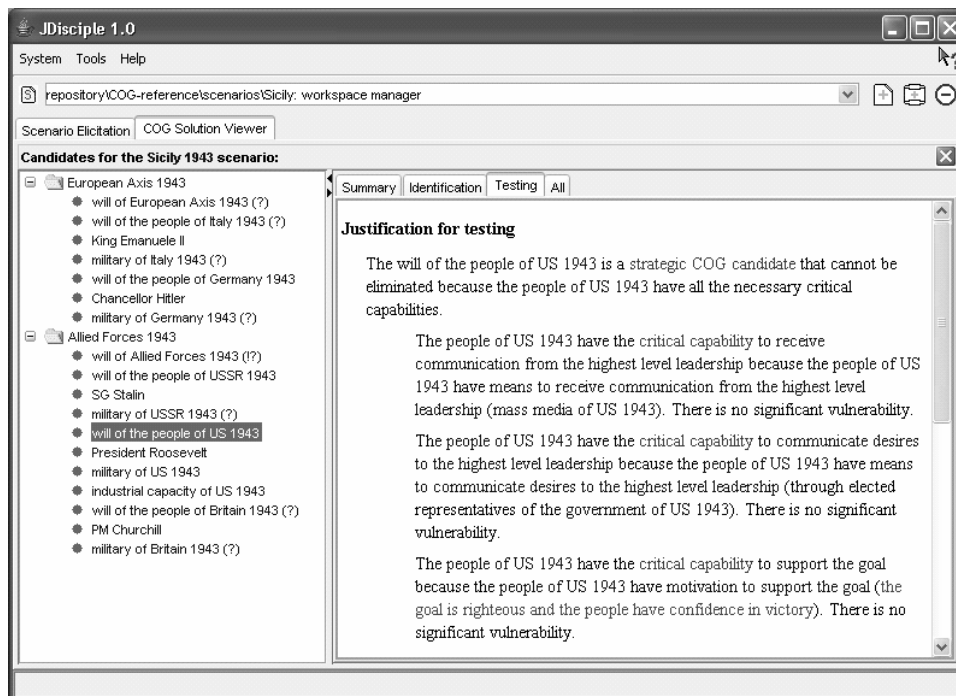


Figure 5: Solution viewer interface.



Figure 6: Training and using the Disciple-COG agent.

Gravity Analysis” course, acted as our main subject matter expert. We have worked with him to teach Disciple-COG in a way that was consistent with how he would teach the students himself (see Figure 6). Therefore he became the main proponent of using Disciple-COG in this course.

Selection of students/end-users: We have invited the students that had a more technical background to attend an information session where we had presented the courses and our research project, asking them to join us in this effort. The goal was to attract those students that were most interested in this different experience.

Involvement of students/end-users: We have involved the students in our research project, making clear that they could contribute to this effort by providing valuable feedback on the scenario description process, on the modeling of the center of gravity analysis process, and on the general system’s characteristics. At the same time, we have presented them recent developments done in response to earlier student feedback. Student feedback was collected in three different ways:

- recording the informal feedback provided during the class use of the system;
- automatic measurements of Disciple use;
- general feedback during the course’s After-Action Review;
- detailed evaluation forms filled-in by students at the end of the course (the evaluation addressed a wide range of issues, ranging from judging Disciple’s usefulness in achieving course’s objectives, to judging its methodological approach to problem solving, and to judging the ease of use and other aspects of various Disciple modules).

Easy to learn and easy to use: A general presentation and a system demonstration was enough for the students to start using the Disciple-COG agent.

Follow established deployment practices: For each new session of the COG course we have developed an improved version of Disciple-COG, based on the lessons learned in the previous sessions. However, we have learned to stop the development of the new version, and to test it as thoroughly as possible, in advance of the course. Then, during the course, despite our desire of continuous

development, we have limited system’s changes to bug corrections.

Immediate support in system’s use: The use of Disciple-COG was supervised by the developers who could provide immediate

assistance in case of software failure. This immediate assistance is of critical importance when using incompletely tested research prototypes. This was a lesson learned from the first time use of the system, in Winter 2001, as discussed below.

Managing user’s expectations: The first version of Disciple-COG was used in Winter 2001. At that time we have provided the system to the students, to use it at home or in class, offering to provide assistance by email. The students expected a commercial-strength system and became frustrated when errors occurred and assistance was not immediately received. As a consequence, the following versions of the system were used in class, under our supervision, any incidents receiving immediate attention. Moreover, we have made it clear to the students that this was an evolving system, and errors are to be expected.

Continuous demonstration of incremental progress: This progress provides confidence to the sponsor and the users to continue to support the research, development and deployment effort. In the case of Disciple-COG, one important way to quantify incremental progress was the evolving size of the knowledge base (see Figure 7), which reflect an increased breath and depth in the modeling of the center of gravity analysis process.

Another way to quantify progress was the students’ satisfaction with the use of the system. For instance, all the 8 students from the Spring 2003 session of the COG course have agreed or strongly agreed with the following statements: “The use of Disciple is an assignment that is well suited to the course’s learning objectives” and “Disciple should be used in future versions of this course”.

Sharing the success with the user: Each COG course

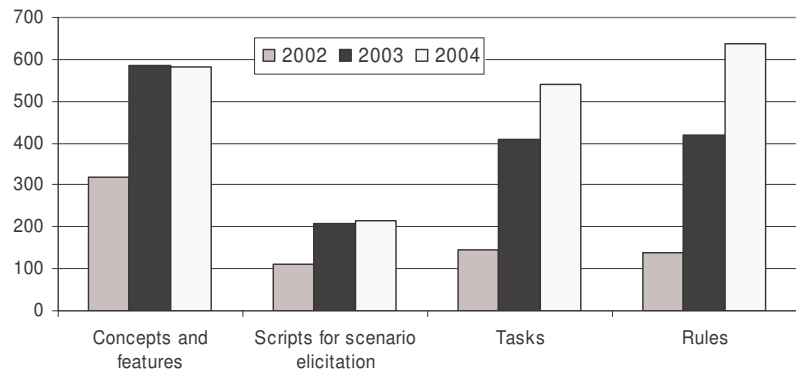


Figure 7: Incremental increase of the knowledge base.

ended with an After Action Review, which was attended by the students, the leadership of the US Army War College, and the representatives from the sponsoring organizations. In addition, we have published joint papers with our collaborators from the US Army War College, both in their and our typical media.

Gradual transition: We are currently in the process of a gradual transition of Disciple-COG, preparing the customer to require less support from us while continuing to use the system.

Availability of funding: GMU has received funding from DARPA's Rapid Knowledge Formation Program and AFOSR's Software and Systems Program that have supported this effort. This allowed GMU to develop and deploy a system at the US Army War College, the beneficiary of this work, without requiring funding from it. Although the US Army War College have latter contributed to the funding of this work, this represented only a small fraction of the necessary funding. Therefore, given the significant funding required for such an effort, and the limited funding possibilities of the potential beneficiaries (senior service colleges, in this case), funding from other sources is an important factor.

We think that it was the combination of all these factors that has led to the success of our deployment. However, the most critical of them are "Working closely with the customer" and "Immediate support in system's use."

Conclusions and Future Plans

Fielding applications is not a common activity for university researchers. The experience reported in this paper shows, however, that it may be very beneficial even for advancing basic research goals. Indeed, our research on knowledge bases and agent development by subject matter experts has benefited from the center of gravity analysis domain which provided a complex, knowledge-intensive, challenge problem. This research has also benefited from its practical application to education. Both the "Case Studies in Center of Gravity Analysis" course and the "Military Applications of Artificial Intelligence" course allowed us to perform thorough experimentations with real experts, resulting in the validation of our methods and providing many ideas for improvements.

On the other hand, the research in center of gravity analysis has benefited from the artificial intelligence research in that the agent development has helped clarify and formalize the center of gravity analysis process, and has actually resulted in its first computational approach. Also, the innovative application of the artificial intelligence and center of gravity research to education, through the use of the Disciple agents, has significantly improved the COG and MAAI courses.

We are currently pursuing two approaches to expand on this success. The first is the plan to develop a new version of Disciple-COG to be used not only at the US

Army War College but also at the other senior service colleges, such as US Marine Corps College and US Air War College. The second plan is to embark on another project with the US Army War College, following the same successful framework of integrating research in artificial intelligence, with research in a specialized domain, and with the development and deployment of prototype systems in education and practice.

Acknowledgments

This research was sponsored by DARPA, AFRL, AFMC, USAF, under agreement number F30602-00-2-0546, by the AFOSR under grant no. F49620-00-1-0072, and by the US Army War College. Several persons supported this effort, including Jerome Comello, William Cleckner, Murray Burke, William Rzepka, Douglass Campbell, David Brooks, David Cammons, and Michael Bowman.

References

- Boicu, M., Tecuci, G., Stanescu, B., Marcu, D., and Cascaval, C.E. 2001. Automatic Knowledge Acquisition from Subject Matter Experts. In *Proceedings of the Thirteenth International Conference on Tools with Artificial Intelligence (ICTAI)*, 69-78. Los Alamitos, CA: IEEE Computer Society.
- Giles, P.K., and Galvin, T.P. 1996. *Center of Gravity: Determination, Analysis and Application*. CSL, U.S. Army War College, PA: Carlisle Barracks.
- Stanescu B., Boicu C., Balan G., Barbulescu M., Boicu M., Tecuci G. 2003. Ontologies for Learning Agents: Problems, Solutions and Directions. In *Proceedings of the IJCAI-03 Workshop on Workshop on Ontologies and Distributed Systems*, 75-82. Acapulco, Mexico, AAAI Press, Menlo Park, CA.
- Strange, J. 1996. *Centers of Gravity & Critical Vulnerabilities: Building on the Clausewitzian Foundation So That We Can All Speak the Same Language*. Quantico, VA: Marine Corps University Foundation.
- Tecuci, G. 1998. *Building Intelligent Agents: An Apprenticeship Multistrategy Learning Theory, Methodology, Tool and Case Studies*. London: Academic Press.
- Tecuci G., Boicu M. 2002. Military Applications of the Disciple Learning Agent. In Jain L, ed., *Advanced Information Systems in Defense and Related Applications*, 337-376. Springer Verlag.
- Tecuci, G., Boicu, M., Marcu, D., Stanescu, B., Boicu, C., Comello, J., Lopez, A., Donlon, J., and Cleckner, W. 2002a. Development and Deployment of a Disciple Agent for Center of Gravity Analysis. In *Proceedings of the Fourteenth Innovative Applications of Artificial Intelligence Conference*, 853-860. Menlo Park, CA: AAAI Press.
- Tecuci G., Boicu, M., Marcu, D., Stanescu, B., Boicu, C., and Comello, J. 2002b. Training and Using Disciple Agents: A Case Study in the Military Center of Gravity Analysis Domain. *AI Magazine* 23(4): 51-68.