Clausewitz Meets Learning Agent Technology

Colonel Antonio M. Lopez, Jr., U.S. Army, Retired;
Colonel Jerome J. Comello, U.S. Army, Retired;
Lieutenant Colonel Michael Bowman, U.S. Army;
Major James J. Donlon, U.S. Army; and
Dr. Gheorghe Tecuci

> *I did not gain a full understanding of many [Clausewitzian] concepts until I had to teach the subject [center of gravity], having read [*On War*] several times.*[1]
>
> —Colonel Huba Wass de Czege[2]

SINCE MICHAEL HOWARD and Peter Paret published their English translation of Carl von Clausewitz's *On War* in 1976, military professionals have been interpreting and finding modern-day meaning in the words of the 19th-century military theoretician.[3] Vice Admiral Stansfield Turner restructured the curriculum of the U.S. Naval War College (USNWC), introducing among many other innovations, the study of Clausewitzian theory. The U.S. Air War College made similar changes in 1978, as did the U.S. Army War College (USAWC) in 1981.[4] Of the many ideas and concepts Clausewitz put forth, his concept of center of gravity (COG) has evoked a significant amount of contention, debate, and writing over the last 20 years. In 1992, a student of the U.S. Army's School of Advanced Military Studies summed up COG's enigmatic nature when he observed, "The concept of center of gravity seems to mean something to everyone, but not the same thing to anyone."[5]

Few writings offer a unified methodology a novice might follow and apply to gain the same wisdom and understanding of the concept that a subject matter expert (SME) has. In 1993, Colonel William Mendel and Colonel Lamar Tooke published an article, titled "Operational Logic: Selecting the Center of Gravity," that provided a means of assessing the validity of an identified COG.[6]

By what logical method can we identify potential strategic COG candidates so we can apply their test of validity? In October 1993, the USAWC Center for Strategic Leadership elicited the knowledge and wisdom of a number of COG experts and developed a methodology for identifying COG candidates and testing their validity. The research evolved into a 1995 master's thesis by Major Timothy J. Keppler and a 1996 monograph by Major Phillip K. Giles and Captain Thomas P. Galvin.[7]

Keppler's thesis specifically explored the question, "Using knowledge engineering techniques, is it possible to distill discernible thought patterns from selected strategists and professional literature to create a useful methodology for applying the center of gravity concept?"[8] His research was an attempt to use systems and knowledge-engineering techniques to model strategic-level thought. The posed research question was answered affirmatively, and a logical methodology was produced to help students and real-world planners consistently apply the COG concept at the strategic and operational levels of war. Keppler's contemporaries at USAWC built on his work and produced the COG monograph that is used as a guide each time the elective course "Case Studies in Center of Gravity Determination" is taught at USAWC.[9]

> ***Joe Strange, of the Marine Corps War College, noted that* On War *is open to a kaleidoscope of individual interpretations when not studied in a professional manner. He suggested that a common language be used. He also recognized that even when groups of people agreed on a common conceptual definition, when the concept was applied to a specific situation, people often identified remarkably different enemy characteristics as the COG.***

The focus in professional literature, however, continued to be on issues of interpretation, confusion, existence, controversy, and utility of the concept, rather than on improving and expanding the USAWC methodology or developing alternative methodologies. In 1996, Joe Strange, of the U.S. Marine Corps (USMC) War College, noted that *On War* is open to a kaleidoscope of individual interpretations when not studied in a professional manner.[10] He suggested that a common language be used. He also recognized that even when groups of people agreed on a common conceptual definition, when the concept was applied to a specific situation, people often identified remarkably different enemy characteristics as the COG.[11]

In 1997, U.S. Navy (USN) Lieutenant Commander Jeffrey Harley wrote that the proliferation of information technology had led to the impression that information is itself a COG, which in turn has confused the role of information and the COG concept.[12] In 1998, USMC Colonel Mark Caucian wrote that centers of gravity just do not exist.[13] In 1999, Republic of Singapore Air Force Major Seow Hiang Lee produced an insightful paper detailing the controversy that still surrounds the COG concept. He suggested four propositions to deal with the confusion as well as three principles on how to use the COG concept.[14] Most recently, Milan Vego, of the USN War College, cited Keppler's research and the COG monograph, but he did not seek to improve or expand the logical methodology therein.[15] By August 2000, USN Commander Jeff Huber had written, "The center-of-gravity theory won't wash if it takes a Zen master decades of rumination from atop the highest peak in Tibet to apply it."[16]

Knowledge Engineering and Learning Agents

Knowledge engineering, a critical activity when developing intelligent agents, is a subfield of artificial intelligence (AI), a branch of computer science. Knowledge engineering is concerned with applying knowledge to solve problems that ordinarily require human expertise. Knowledge engineers perform the following three major functions:

1. Identify problem domains.
2. Perform knowledge acquisition to understand how SMEs solve problems and to elicit their problem-solving knowledge.
3. Construct intelligent agents that incorporate the problem-solving knowledge acquired from SMEs.

Knowledge engineers identify domains that give an organization a significant payoff in cost savings or in providing an advantage over a competitor, if the organization can apply automated knowledge to problems encountered. Appropriate problem domains for knowledge engineering are domains where humans solve problems that are unstructured; have a large number of variables, some of which have unknown values because of incomplete information; have multiple or conflicting goals; and make use of highly specialized knowledge. The COG concept is certainly an appropriate domain for knowledge engineering.

A necessary condition for successful knowledge acquisition is access to experts who can solve the problem well and who know how to communicate or who can demonstrate that expertise. Keppler realized USAWC military professionals were recognized as being able to apply the COG concept effectively and consistently.[17] Keppler's method of knowledge acquisition was to interview SMEs, observe practical exercises, and elicit knowledge from then current professional literature, which is the traditional knowledge-acquisition approach that requires significant interaction between trained knowledge engineers and SMEs.

Successful knowledge acquisition contributes to the development of an intelligent agent. An intelligent agent is a computer program that perceives its environment, interprets perceptions, draws inferences, solves problems, determines actions, and acts on its environment to realize a set of goals for which it was designed.[18] By 1995, the USAWC had distilled the acquired knowledge into a methodology for COG determination. While the ultimate goal was to build an intelligent agent based on this knowledge and its resulting methodology, as an interim solution, the USAWC developed a decision-support system to guide users through the COG-determination process and related considerations. The software was used to facilitate the COG course until the end of the 1998 academic year.

One of the primary impediments to learning-agent construction at the USAWC was the time, effort, and expertise needed to formalize the acquired knowledge and to develop an agent. In the traditional knowledge-acquisition approach, knowledge engineering involves transferring and transforming an SME's knowledge into a form usable by an intelligent agent. A skilled knowledge engineer ordinarily performs this highly technical process, which is time consuming, error prone, and inefficient. An alternative approach is to use a computer-based learning agent, which can acquire and maintain the SME's knowledge with only limited assistance from a knowledge engineer.[19]

The Learning Agents Laboratory (LALAB) at George Mason University (GMU) developed the new approach, calling it Disciple. Disciple is an apprenticeship, multi-strategy learning approach for developing intelligent agents. An SME teaches a Disciple agent (software programs that run on common laptop or desktop computers) how to perform domain-specific tasks in a similar manner in which the SME would teach an apprentice—by giving examples and explanations and by supervising and correcting behavior.[20]

The Disciple approach has been successful in a number of different applications, including assessment, planning, design, and critiquing tasks. A recent successful military application of Disciple involved critiquing courses of action for tactical military plans.[21] Agent technology, combined with the

In the traditional knowledge-acquisition approach, knowledge engineering involves transferring and transforming an SME's knowledge into a form usable by an intelligent agent. A skilled knowledge engineer ordinarily performs this highly technical process, which is time consuming, error prone, and inefficient. An alternative approach is to use a computer-based learning agent, which can acquire and maintain the SME's knowledge with only limited assistance from a knowledge engineer.

continued professional interest in strategic and operational COG determination, presented an excellent opportunity to advance the knowledge-acquisition work of Keppler, Giles, and Galvin to develop an intelligent agent. Therefore, the USAWC and GMU are developing Disciple-COG, an intelligent agent for COG determination that can be taught directly by SMEs, with limited assistance from knowledge engineers. This work is a collaborative effort between the Department of Military Strategy Planning and Operations and the Center for Strategic Leadership, at USAWC, and LALAB, at GMU.

Developing the Theory

Clausewitz was a theoretician who attempted in his books, essays, and notes to lay out a system of thought regarding war. His writings, although not completely satisfactory to himself at the time of his death, present a theoretical model based on reason and logic against which judgments can be made about real phenomenon.[22] Like the initial models theoreticians presented in other disciplines (mathematics, for example), good models deserve further development so they can become better models. The calculus of Gottfried Liebnitz and Isaac Newton is not that being taught today to engineering students at the U.S. Military Academy. The present calculus model is the result of logic and reason refined over many years. Today's engineering students are using graphing calculators and computer algebra systems to demonstrate and calculate solutions to problems that Liebnitz and Newton would not have dreamed of attempting. So too are military professionals and knowledge engineers called on to continue developing Clausewitz's COG theory, refining it with logic and reason and incorporating the latest technology to analyze more difficult scenarios than those of Clausewitz's day.

In a classic work, Douglas B. Lenat and Edward A. Fiegenbaum stated their Empirical Inquiry Hypothesis, which claims that the best action AI researchers can take to further the development of the field is to take their ideas, incorporate them into programs, run the programs, and see where they fail.[23] This is where AI researchers will learn the most. The same can be said for military theoreticians. They need to take a theory, such as Clausewitz's COG; incorporate it into Disciple-COG; teach Disciple-COG to determine and analyze strategic and operational COGs; and see what Disciple-COG does not do well. Doing this will help them gain greater insight into the theory as well as to refine a methodology for its understanding and application by students. Ultimately, Disciple-COG will become an intelligent partner in applying the theory to present-day scenarios.

In continuing to develop COG theory, we are making use of the various historical case studies prepared by USAWC faculty and students. In the 1998 *Military Review* article "Center of Gravity and Strategic Planning," Steven Metz and Frederick Downey caution, "While individual historical studies are useful for a strategic planner, their value is eroded by the absence of any general guidelines or conclusions collated from a number of cases."[24] We agree fully, and our approach abstracts such general guidelines from the cases studied. Disciple-COG will learn from examples, explanations, analogy, and its own experimentation based on a wealth of individual historical case studies that experts and students provide. Disciple-COG will synthesize these cases to learn principles that are generally applicable, without having been explicitly told them. When required to do so by a student or SME, Disciple-COG will explain in detail the reasoning it used to draw its conclusions. This reasoning might be based on one specific historical scenario that serves as an analogy for the present problem, or it might be based on fragments of knowledge from many different historical scenarios that, when recalled and reconfigured under the present problem, give a plausible solution. If Disciple-COG cannot use its historical knowledge to solve a given problem, it will seek further guidance and training from a COG SME, by which it will further improve its knowledge and expertise.

Transforming any theory into something an intelligent agent can understand cannot be accomplished overnight. As a start, we can draw on the work of Keppler, Giles, and Galvin and the SMEs their works cite. Also, however, this process needs technical expertise not always readily available. The USAWC is an ideal environment for continued access to

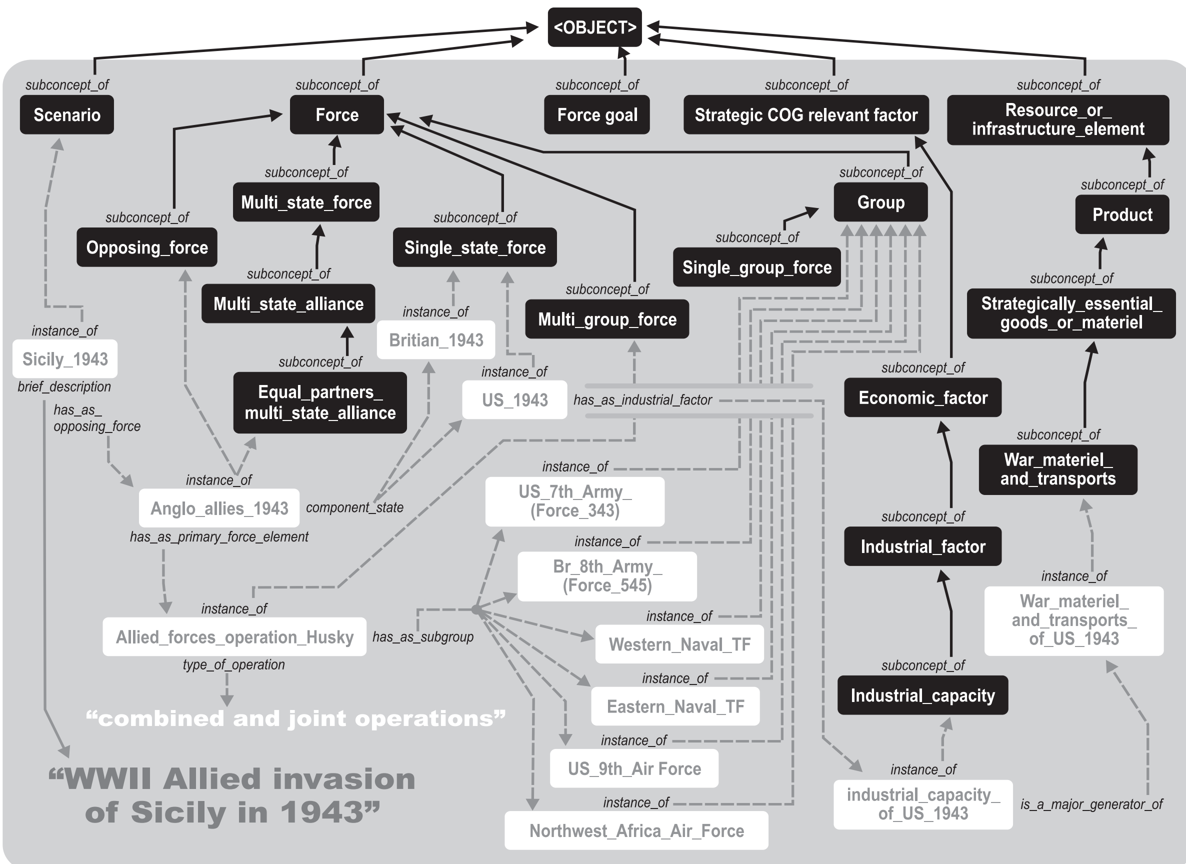

Figure 1. A fragment of the object ontology developed for Disciple-COG.

current thought on COG. In fall 2000, Murray Burke, Defense Advanced Research Projects Agency (DARPA) program manager for the Rapid Knowledge Formation program, directed LALAB to partner with the USAWC to develop Disciple-COG to advance the state of the art in conducting knowledge acquisition from domain experts.[25] DARPA, the Air Force Office of Scientific Research, and the Air Force Research Laboratory of the Air Force Materiel Command (AFMC) are funding the LALAB's research. In the first year of this effort, LALAB and USAWC have had considerable success in acquiring domain knowledge and have begun agent development based on this knowledge. We describe below how Disciple was successfully integrated into a course in COG determination to elicit knowledge based on historical cases.

Developing the Agent

The main phases of agent development include the following:

- Customizing the agent shell.
- Developing the agent's object ontology.
- Modeling the problem-solving process.
- Teaching the agent.
- Verifying and validating the agent.

In the first year of Disciple-COG development, the focus has been on identifying strategic COG and has touched on all phases of agent development. In general, customizing the Disciple shell for a particular application consists of developing new modules or, at least, extending and adapting existing modules, to satisfy the current application's requirements. The object ontology consists of specifying the objects and type of objects, with their properties and relationships, from the application domain. For the COG domain, developing the object ontology was based on the previous works of Keppler, Giles, and Galvin, and on the detailed analysis of two case studies provided by the USAWC professor who taught the COG course in January 2001. The two case studies were the Sicily and Okinawa campaigns of World War II. Over 100 pages of diagrams document the developed ontology. Figure 1 presents a small fragment of this ontology with selected instances from the Sicily scenario.

The object ontology represents everything the agent "knows" about the subject at hand. Figure 1 contains abstract concepts, depicted in black, as well as specific instances of those concepts, depicted in dark gray. For example, as defined in Disciple-COG, industrial capacity is an abstract concept, and the

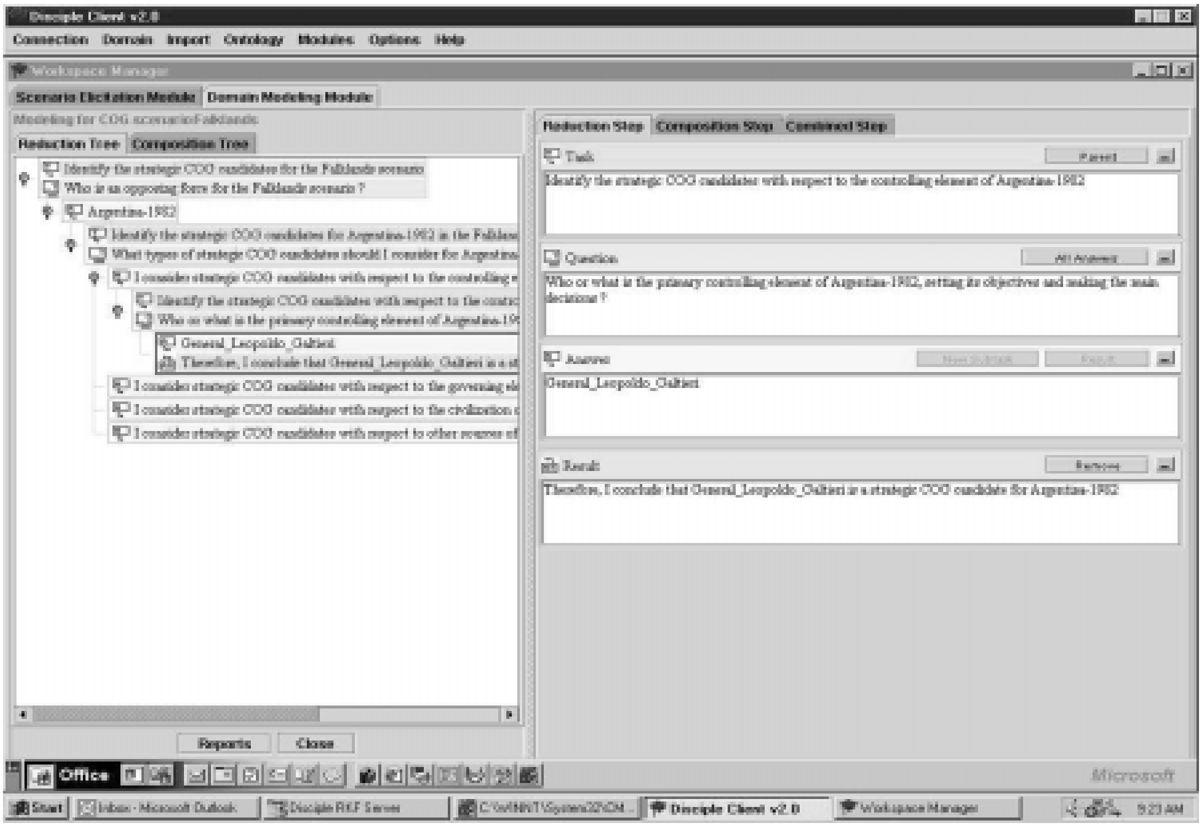

Figure 2. Scenario Elicitation Tool of Disciple-COG with information concerning the Falklands scenario.

specific industrial capacity of the United States in 1943 is an instance of that concept.

In addition to the definitions of abstract concepts and specific instances of those concepts, the ontology represents other important relationships between concepts. Arrows illustrate the presence of a relationship, and light gray labels specify their name. Many abstract concepts are taken from the COG monograph and will not vary from scenario to scenario. The specific instances in a knowledge base, however, are features that describe a particular scenario. When developing the initial ontology for Disciple-COG based on Sicily and Okinawa, knowledge engineers studied the historic cases and added the instances and relationships needed to describe them to the agent.

Figure 1 illustrates only a portion of the complete ontology resulting from this initial knowledge-engineering effort. For example, the only instance of an opposing force shown in Figure 1 is the Anglo_allies_1943; however, the European_Axis_1943 force is also present as an opposing force in the complete ontology, and it, in turn, has component states (Italy_1943 and Germany_1943) as well as primary force elements (Axis_forces_Sicily). Similarly, while one strategic COG relevant factor (the economic factor) is depicted, additional factors found in the COG monograph (psychosocial, political, historical factors, and so on) are found in the complete ontology.

Although knowledge engineers prepared the Sicily and Okinawa scenarios to give LALAB researchers some typical concepts, relationships, and instances with which to customize the Disciple agent shell to Disciple-COG, it was USAWC students who, in the spring of 2001, validated the usability of the initial ontology and expanded on it. They accomplished this during the elective COG course.

The students used the Scenario Elicitation Tool, a new, customized Disciple-COG component, to describe their scenarios by answering multiple-choice questions derived from the agent's ontology and elaborating on those answers with descriptions in unrestricted English. The Scenario Elicitation Tool directly supported knowledge-base development by eliciting key instances and relationships from students and linking them to the initial ontology. The USAWC students thus developed scenarios about the following historical case studies: Malaya 1941-42, Leyte 1944, Inchon 1950, Vietnam 1968-75, Falklands 1982, Grenada 1983, Panama 1989, and Somalia 1992-94.

Figure 2 shows the Scenario Elicitation Tool displaying selected entries from the Falklands scenario. The left-hand side of the display shows the table of contents created for each opposing force entered and

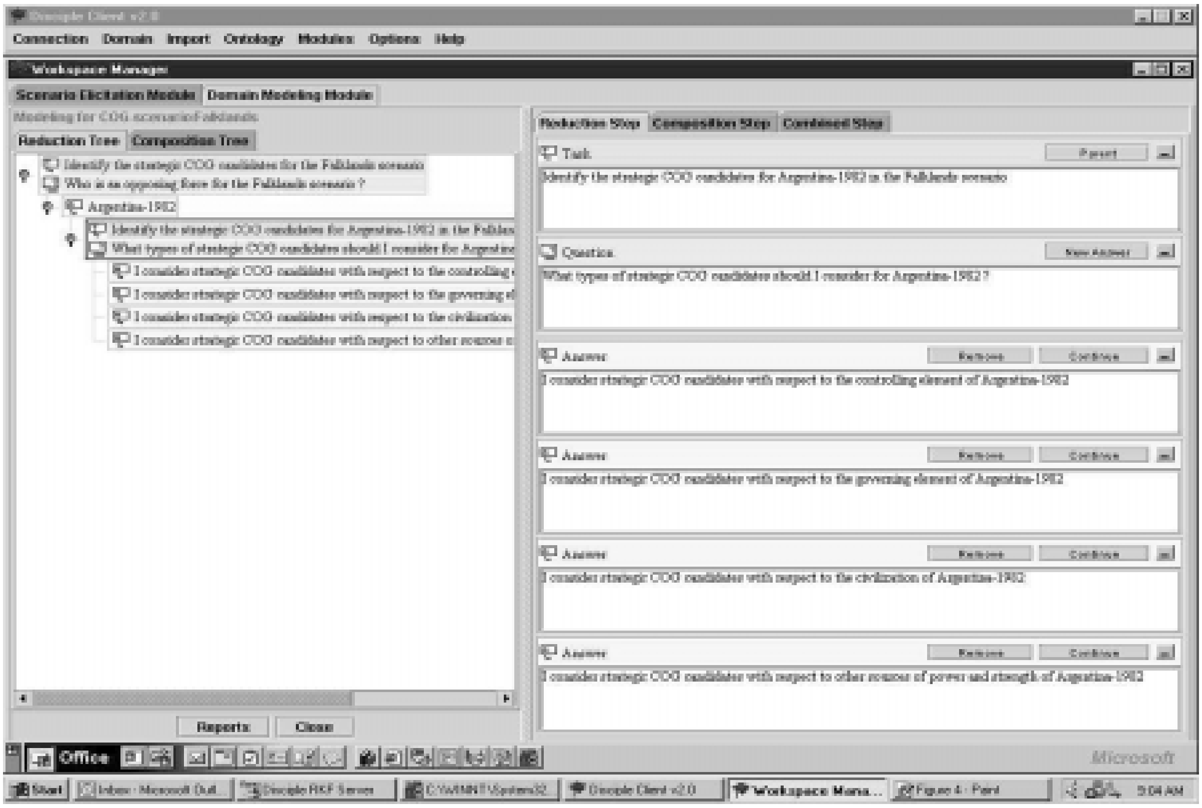

Figure 3. Employing a task-reduction methodology using the Domain Modeling Tool.

the identified strategic and operational COG candidates. Only the Argentina-1982 table of contents is clearly visible in Figure 2, but the vertical slide bar can be used in Disciple-COG to reveal the table of contents for Britain-1982.

Using the Scenario Elicitation Tool, the student highlights a topic in the table of contents (Falklands is highlighted in Figure 2) and enters information for that topic to the right of the table of contents. The right portion of Figure 2 shows that a student entered the scenario name, a subject summary, a brief description, and the opposing forces involved. By entering specific information for Argentina-1982 under *Composition of forces* the student caused *Cooperation_between_members_of_Argentinean_Armed_Forces* to appear under the folder *Strategic COG candidates*. Likewise, information entered under *Control and governing elements,* supported by facts stated in *Historical factors, Military factors,* and *Political factors,* produced strategic COG candidates *General_Leopoldo_Galtieri* and *Military_Junta*. The *Argentinean_Unions* became a strategic COG candidate because of the information entered under *Civilization* and supported by facts in *Economic factors*. In the table of contents for Figure 2, some *Operational COG candidates* are visible. Since the Scenario Elicitation Tool also organizes and formats the report that the students were required to produce for the COG course, this component had to be made available even though it is not the focus of the research for the first year.

Part of the debate over the COG concept is whether it can be applied successfully to operations other than war. To support further study in the area, USAWC students created Disciple-COG scenarios for U.S. operations in Panama in 1989 and Somalia during the 1990s. Because of the clans involved in the Somalia 1992-94 scenario, the initial ontology was expanded beyond the World War II concepts found in the Sicily and Okinawa scenarios. Additional concepts such as *Chief_and_tribal_ council* and *Democratic_council_or_board* were needed to develop the possible types of a governing body for a clan or a tribe. Similarly, the Panama 1989 scenario caused ontology expansion to include concepts such as *drug_cartel* and *crime_family*. Ontology expansion can be expected to continue with each new scenario USAWC students visit during the COG course.

The Scenario Elicitation Tool develops the ontology and captures instances and relationships, but it does not enable autonomous reasoning by Disciple-

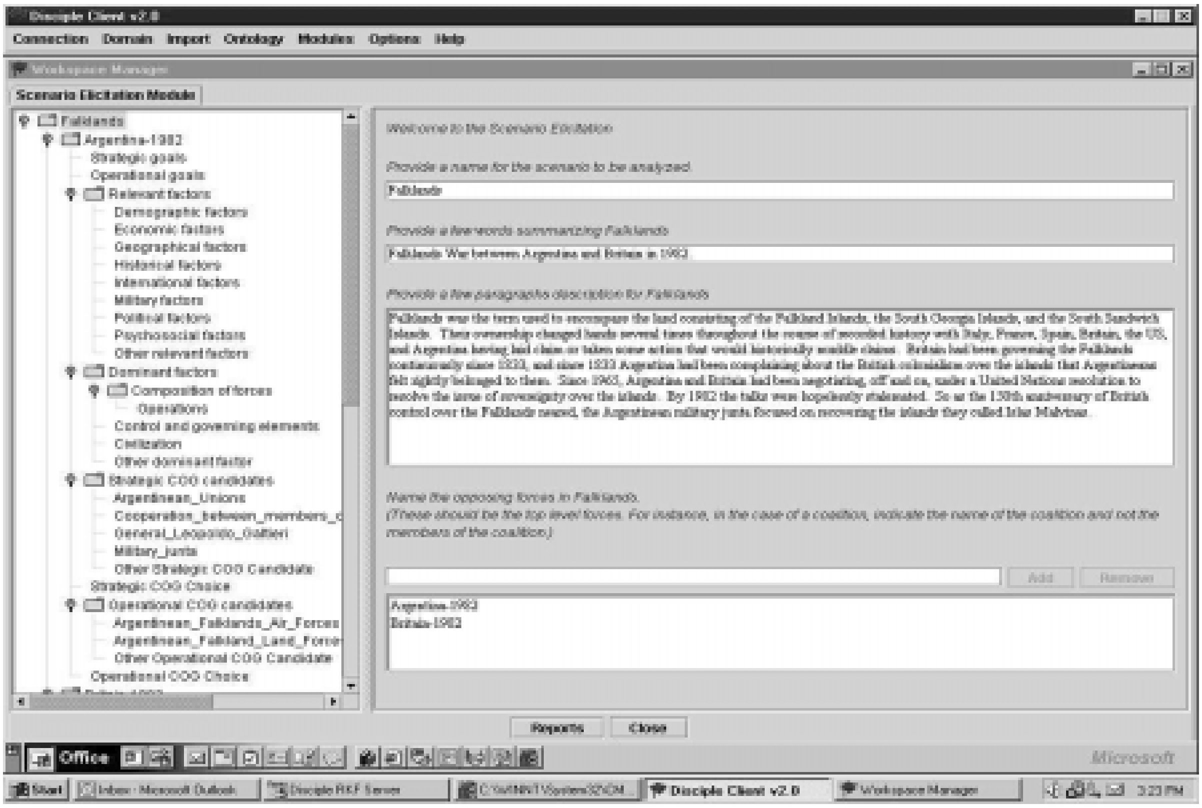

Figure 4. A task-reduction thread that ends with a result.

COG. For this we need to model the problem-solving process, which is the next step in agent development. Another Disciple component, the Domain Modeling Tool, supports this modeling. The Domain Modeling Tool is based on a task-reduction methodology that allows the user to state a task, ask a question about the task, and provide one or more answers. The fundamental concept at work in the Domain Modeling Tool is that a complex problem can be successively reduced to simpler subproblems until the subproblems are simple enough to be solved immediately. The solutions to the subproblems can then be successively combined to produce the solution to the initial problem. This general concept has been given many names including problem or task decomposition, factorization, and task reduction. We prefer the term task reduction.

Figure 3 shows the Domain Modeling Tool as a student has begun doing task reduction in the Falklands scenario. The first task-reduction step includes the task "Identify the strategic COG candidates for the Falklands scenario"; the question "Who is an opposing force for the Falklands scenario?"; and the answer "Argentina-1982." The answer suggests to the student a subtask (the current task) and further task reduction. The follow-on question from the current task produces four answers, taken from the COG monograph, which the student must analyze further.

During task reduction, a task is successively reduced to simpler and simpler tasks. Each subsequent reduction step is based on the consideration of some relevant factors, expressed as a question. Each answer to a question guides the user to reduce the current task to a simpler one. Eventually each task-reduction sequence terminates with a result.

Figure 4 shows a completed pattern of reasoning that identifies a strategic COG candidate. The right portion of Figure 4 illustrates a task-reduction sequence; the left portion shows where this trend of thought exists in the overall problem-solving scheme being developed.

Modeling the problem-solving process of strategic COG identification was the most difficult and time-consuming aspect of the work done in the COG course. The students, not SMEs in COG determination, selected and completed at least two task-reduction sequences that identified candidate COGs for their scenarios. Some students performed additional task reductions. The thought patterns they derived were used in further agent development to teach Disciple-COG how to identify strategic COG candidates.

Summary

In the 19th century, Clausewitz presented a theory about war that was rediscovered after the Vietnam war. Since then, several professionals knowledgeable in the art of war have given their interpretations of Clausewitz's COG theory. In 1996, after years of effort by knowledge engineers to acquire and synthesize knowledge from COG SMEs, USAWC published the COG monograph, a methodology for COG determination, analysis, and application.

The COG monograph and several historical case studies formed the basis for continued development of the COG theory by teaching it to an intelligent agent called Disciple-COG. Through a partnership between USAWC, LALAB, and DARPA, students in the USAWC COG course used LALAB's Disciple-COG to develop their assigned historical scenarios and to model the way they identified their strategic COG candidates. Significant strides have been made unifying COG theory and learning-agent technology in Disciple-COG. ***MR***

Successful knowledge acquisition contributes to the development of an intelligent agent. An intelligent agent is a computer program that perceives its environment, interprets perceptions, draws inferences, solves problems, determines actions, and acts on its environment to realize a set of goals for which it was designed.

NOTES

1. Carl von Clausewitz, *On War*, trans. and ed. Michael Howard and Peter Paret (New Jersey: Princeton University Press, 1976).
2. Huba Wass de Czege, "Clausewitz: Historical Theories Remain Sound Compass References: The Catch is Staying on Course," *Army* (September 1988): 38. Wass de Czege continues to be a visionary in military affairs. In September 2000, he participated in the Objective Force Combat Service Support Workshop at the USAW.
3. Clausewitz.
4. John B. Saxman, *The Concept of Center of Gravity: Does It Have Utility in Joint Doctrine and Campaign Planning?* (Monograph, School of Advanced Military Studies, Fort Leavenworth, KS, 1992), 1.
5. Ibid., 4.
6. William W. Mendel and Lamar Tooke, "Operational Logic: Selecting the Center of Gravity," *Military Review* (June 1993): 5.
7. Timothy J. Keppler, "The Center of Gravity Concept: A Knowledge Engineering Approach to Improving Understanding and Application" (Master's thesis, U.S. Army Command and General Staff College, Fort Leavenworth, 1995, available through Defense Technical Information Center, Fort Belvoir, VA); Phillip K. Giles and Thomas P. Galvin, *Center of Gravity: Determination, Analysis, and Application* (Carlisle Barracks, PA: Center for Strategic Leadership, USAWC, 1996), online at <http://carlisle-www.army.mil/usacsl/publications/gravity.pdf>. The process-flow diagram that accompanies this monograph is available from Jerome J. Comello.
8. Keppler, 1.
9. Comello recently offered Course 319jw (joint warfighting), "Case Studies in Center of Gravity Determination," in terms II and III of the 2000-01 academic year.
10. Joe Strange, *Centers of Gravity & Critical Vulnerabilities: Building on the Clausewitzian Foundation So That We Can All Speak the Same Language* (Quantico, VA: U.S. Marine Corps University Foundation, 1996), 26.
11. Strange, 31.
12. Jeffrey Harley, "Information, Technology, and the Center of Gravity," *Naval War College Review* (Winter 1997), online at <www.nwc.mil/press/Review/1997/winter/art4wi97.htm>.
13. Mark Cancian, "Centers of Gravity Area a Myth," *The U.S. Naval Institute Proceedings* (September 1998), online at <www.usni.org/Proceedings/PROcancian.htm>. This article was a Colin L. Powell Joint Warfighting Essay Contest Winner.
14. Seow Hiang Lee, Republic of Singapore Air Force, *Center of Gravity or Center of Confusion: Understanding the Mystique* (Maxwell Air Force Base, AL: Air Command and Staff College, 1999), online at <www.au.af.mil/au/database/research/ay1999/acsc/99-110.htm>.
15. Milan Vego, "Center of Gravity," *Military Review* (March-April 2000), online at <www-cgsc.army.mil/milrev/English/MarApr00/veg.htm>.
16. Jeff Huber, "You Can't Defy the Laws of Gravity," *The U.S. Naval Institute Proceedings* (August 2000): 38.
17. Keppler, 7.
18. *See* Gheorghe Tecuci, *Building Intelligent Agents* (San Diego, CA: Academic Press, 1998), 1.
19. Ibid., 2.
20. Ibid., 13.
21. *See* Michael Bowman, Gheorghe Tecuci, and Mihai Boicu, "Intelligent Agents in the Command Post," *Military Review* (March-April 2001): 52-53.
22. Wass de Czege, 38.
23. Douglas Lenat and Edward Feigenbaum, "On the Threshold of Knowledge," *Artificial Intelligence* (1991), 204. Feigenbaum, who has served as chief scientist for the U.S. Air Force (USAF) since 1994, received the USAF Exceptional Civilian Service Award.
24. Steven Metz and Frederick Downey, "Center of Gravity and Strategic Planning," *Military Review* (April 1998): 23.
25. Murray Burke, Rapid Knowledge Formation Program Description (January 2000), online at <http://dtsn.darpa.mil/iso/programtemp.asp?mode=331>.

Colonel Antonio M. Lopez, Jr., U.S. Army Reserve, Retired, holds the Chair of Artificial Intelligence in the Knowledge Engineering Group of the Center for Strategic Leadership at USAWC. He also holds the Conrad N. Hilton Endowed Chair in Computer Science at Xavier University of Louisiana. He served as the assistant chief of staff G6, G3, and chief of staff with the 377th Theater Support Command.

Colonel Jerome J. Comello, U.S. Army, Retired, is professor of Military Studies at USAWC. He was professor of National Security Studies at the U.S. Marine Corps (USMC) Command and Staff College and professor of Strategy Plans and Operations at the Center for Strategic Leadership. He graduated from USMC, the College of Naval Command and Staff, and the National War College. He received an M.S. from George Washington University and a doctorate from Temple University.

Lieutenant Colonel Michael Bowman is a graduate of USAWC and is a doctorate candidate at George Mason University. He received a B.S. from Ouachita Baptist University and an M.S. from the Naval Postgraduate School. He was the Army Product Manager for Communications and Intelligence Support Systems, and he has had a variety of acquisition, automation, and tactical assignments at the U.S. Defense Intelligence Agency and the U.S. Military Academy (USMA).

Gheorghe Tecuci is professor of Computer Science and director of the Learning Agents Laboratory at George Mason University, and he is a member of the Romanian Academy. He received an M.S. and doctorate degrees from the Polytechnic University of Bucharest and the University of Paris-South. He has published more than 100 scientific papers and five books, most of which have been in the artificial intelligence (AI) areas of intelligent agents, machine learning, and knowledge acquisition.

Major James J. Donlon is the director of the Knowledge Engineering Group at USAWC. He received a B.S. from the University of Delaware and an M.S. from Northwestern University. He conducts applications engineering, education, and applied AI research as a U.S. Army systems engineer.