

# MODELING TERRORISTS

NEW SIMULATORS COULD HELP INTELLIGENCE  
ANALYSTS THINK LIKE THE ENEMY  
**BY HARRY GOLDSTEIN**

“There are tools where they build a world in a bottle. They put down every single mosque, river, camel, and school in, say, Saudi Arabia. Then they have millions of software agents who each have desires, grievances, all these different variables. They go about their little lives and then you ask a question: What if we build a McDonald’s in Mecca? Does this lead to more people joining terrorist groups or not?”

—Gary Ackerman,  
*Director of the Center for Terrorism  
and Intelligence Studies*



**BARRY SILVERMAN**  
peers deep into the  
heart of darkness to  
find what makes  
terrorists tick.

BARRY SILVERMAN pecks at the keyboard, and suddenly his computer monitor is showing him the view down a scary-looking alley in the Bakhara market in Mogadishu, Somalia. On the big screen, Silverman sees the market through the eyes of his avatar, a software soldier. It's a detailed scene, on a par with what you'd see in today's best first-person shooter video games: in the market's narrow lanes, militiamen scurry about, checkered headdresses flapping. It has rained recently, and the gray masonry walls of buildings surrounding the market are water stained. The streets are empty except for some abandoned cars and the smoldering wreckage of two helicopters. Silverman's cybertrooper is part of a virtual squad replaying the scenario described famously in Mark Bowden's 1999 best seller, *Black Hawk Down*, in which U.S. Army Rangers attempted a rescue after fighters loyal to warlord Mohamed Farrah Aidid shot down two U.S. UH-60 choppers.

The Ranger that Silverman controls wanders only a few steps toward the downed helicopters before he encounters a suicide bomber who blows them both to bits.

Silverman, an electrical and systems engineering professor at the University of Pennsylvania, in Philadelphia, restarts the simulation. As his Ranger avatar scans the scene, Silverman describes the attributes of each character—or synthetic human agent—he encounters. He knows them all intimately, their motives, emotions, and physiologies, as well as their political, religious, and moral leanings. He should; he and his group created every last one of them.

Through the Ranger's gunsight we see a Somali woman dressed in flowing blue robes and matching head scarf walking with a militiaman clad in an ankle-length white garment. Raising his voice above sporadic gunfire and the crunch of boots, Silverman is explaining that some of his graduate students spent an entire semester studying the behavior of Somali women and their value systems.

BILL CRAMER/WONDERFUL MACHINE

He points to the screen as the woman allows the man to hold her in front of him. “This is not scripted,” he says. “Somali women will act as shields for their men....She is acting according to her values, her physiology, her stress, which are tuned to a person in that culture, and she of her own volition does the things that you see unfold here.”

Silverman, whose sleepy brown eyes and deliberate speech belie a dry wit, gets the man in the crosshairs of his Ranger’s gun-sight. “He’s already upset, because we’ve been over there trying to kidnap the whole leadership of his tribe for a while now. We’re not as innocent as I’m playing here; I’m already sort of labeled....” Gunshots ring out, bullet casings clink on the ground. “They’re looting...and now I’m trying to chase them away.” Suddenly, chaos. An explosion rocks the market, followed by a spray of gunfire. “He’s shooting back at me, and it’s hard for me to aim at him because he’s got the woman there”—pop, pop, and then a moan as Silverman drops the militiaman. “Oh,” he says, surprised by his own marksmanship. “I got him.”

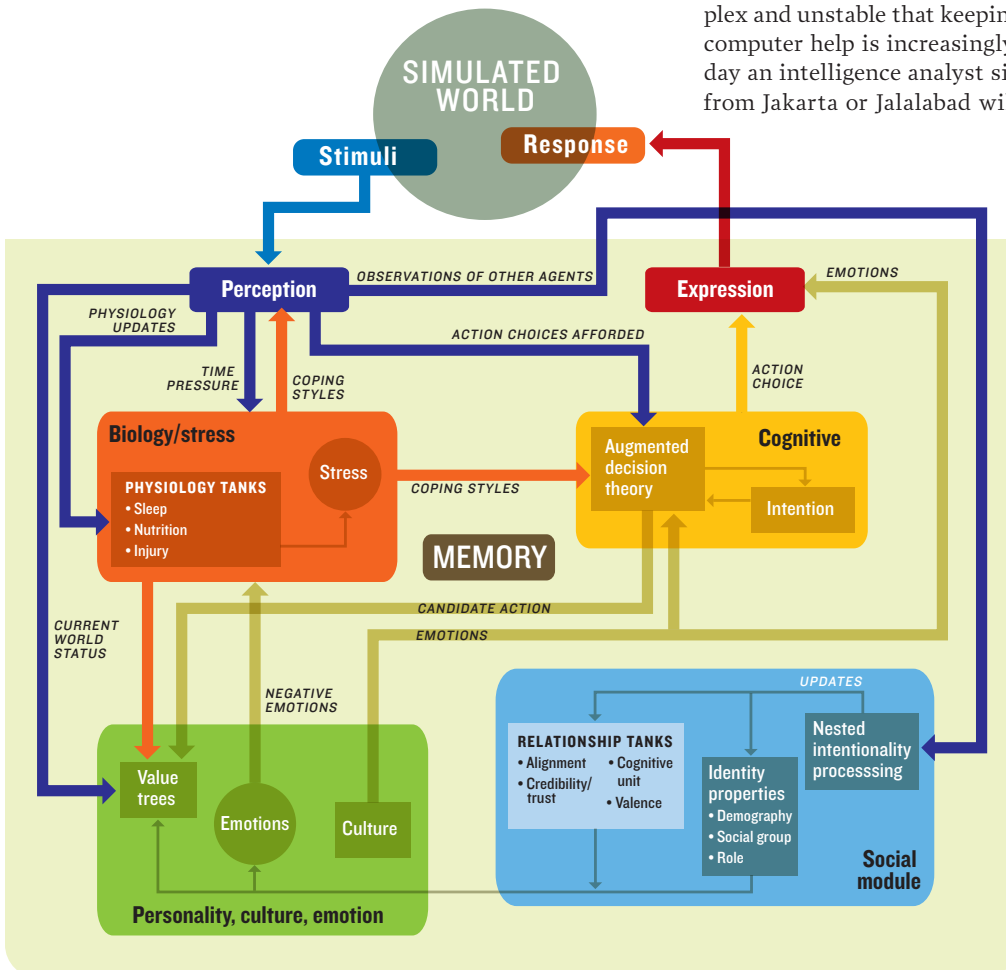
The woman slinks away. “She’s now leaving, because she has no reason to obey him anymore. He’s dead.”

The mere fact that his agents improvise based on individualized sets of complex rules instead of acting according to a rigid script would be enough to make Silverman a rock star among game developers. In fact, the Bakhara market simulation looks like a first-person shooter because it’s based on the *Unreal Tournament* game engine from Epic Games, which renders the scenes and drives the interactive environment. Nonetheless,

though his characters are brought to life by a commercial game engine, the software that gives his characters their individual identities is generations beyond anything you’ll find in a video game today. Silverman’s agents, along with those being developed by a few other teams, are about the closest a computer comes to simulating the thought processes of a real person. Similar work is being done by *The Sims* cocreator John Hiles at the Naval Postgraduate School, in Monterey, Calif., and Jonathan Gratch and Stacy Marsella at the University of Southern California’s Institute for Creative Technologies, in Marina Del Ray.

“This is really at the cutting edge of computational behavior modeling,” says Gary Ackerman, who, as director of the Center for Terrorism and Intelligence Studies, a think tank in San Jose, evaluates these kinds of programs for various U.S. government agencies. “They have been more successful than I thought they could have been, pretty much frighteningly so at some points.”

Silverman’s group is one of several driven by a hypothesis that seemed far-fetched even five years ago: that computers equipped with the right software can give vital insights into the minds and motives of terrorists and the structure and critical links in their organizations. The work is part of a larger effort, much intensified after 9/11, in which the U.S. intelligence community, in particular, is looking for better ways to identify terrorists, determine their capabilities, and predict where and when they will strike. Different forms of the software are aimed at military officials, who are already using such programs to train officers and troops, and at intelligence analysts, who are finding that the shadowy, shifting organizations they must study are so complex and unstable that keeping track of all the variables without computer help is increasingly unrealistic. The hope is that one day an intelligence analyst sitting at a desk thousands of miles from Jakarta or Jalalabad will be able to make preternaturally



**GETTING INSIDE THE MIND:** Barry Silverman’s team at the University of Pennsylvania models the individual terrorist. Each of Silverman’s simulated terrorists responds to stimuli through a complex process that mimics human thinking. An agent’s perception is influenced by various physical and psychological stressors and coping styles. That perception is then viewed in light of an agent’s values, emotions, and culture. Once the agent has processed its perception of the stimuli, it must decide how to best attain its goals. It evaluates which actions will advance its values, but this is subject to constraints of its social relations, emotions, and stress. When the agent determines what it wants to do, it expresses that decision through action.

SOURCE: BARRY SILVERMAN

good guesses about who is likely to commit violent acts, and to advise policy-makers on specific ways to prevent an attack.

Silverman's group focuses on individual agents, but other modelers take a more organizational approach, simulating large-scale social networks on supercomputers and churning out trillions of bytes of data. Models built by Edward MacKerrow at Los Alamos National Laboratory, Charles Macal at Argonne National Laboratory, Alok R. Chaturvedi at Purdue University, Desmond Saunders-Newton at BAE Systems, and Kathleen Carley at Carnegie Mellon University use thousands or millions of relatively simple agents to examine how networks form and mutate, how individuals communicate, and who leads and who follows. Carley's programs, which process real data, stand out for their ability to help analysts imagine how a terrorist network might adapt—or not—after its leader is killed or captured.

Such work, concentrated in the United States and sustained by tens if not hundreds of millions of dollars in funding by various intelligence organizations, including the CIA and the Defense Intelligence Agency, points to a new era in training and intelligence analysis. The experts developing these systems are reticent about exactly how their programs are being used. But outside observers say it is a good bet that software designed to identify the critical people in a terrorist organization will be used—if it hasn't been already—to draw up lists that prioritize which people should be killed or captured so as to do maximum damage to the organization.

That worries some experts, who caution that even when the models are fed by the best available intelligence, they should never be trusted to determine, by themselves, whether someone should live or die. "A simulation is by its nature speculative, and you don't go out and kill people based on speculation," says Steven Aftergood, director of the Project on Government Secrecy for the Federation of American Scientists, in Washington, D.C.

Many modelers emphasize that such simulations are not intended to replace analysts but to augment their abilities to ferret out key individuals, break up covert cells, and prevent the kinds of surprises that lead to devastating terrorist successes. That still leaves one huge question unanswered, skeptical insiders say: Will analysts, many of whom struggle just to stay abreast of the information they are inundated with every day, bother to use these modeling tools if they ever become widely available?

**INTELLIGENCE IS BY ITS VERY NATURE** hazy and fragmentary. Its practitioners' successes must remain secret, while their worst failures erupt in near-real time for all the world to see. For U.S. Intelligence, the attack on Pearl Harbor, the collapse of the Soviet Union, and 9/11 will resound indefinitely. Yet, in all three of those misses, scraps of information collected before the events hinted at what was to come, only to languish undigested or even unnoticed by analysts.

Part of the problem is the way analysts work, which predisposes them to what the 9/11 Commission termed "failure of imagination." Analysts are experts, with advanced degrees in areas like economics or German literature or social psychology, who know one country or group or industry extremely well. For many, the only things that diverted them from careers in academia were patriotic inclinations or the quiet thrill of poring over deciphered intercepts, satellite photos, and data gathered by spies.

This academic culture flourished during the Cold War. Back then, analysts spent much of their time weighing pieces of classified information and thinking about strategies to achieve long-term policy goals. For the vast majority of analysts, anticipating attacks on the homeland wasn't in the job description. But after 9/11, two developments combined to make life for many analysts much more hectic. One was the urgent need to more closely track elusive enemies who were obviously committed to killing people and destroying property. The other was the establishment of the Internet as the primary source of publicly available information—and the preferred means of terrorist communication. The Internet hugely increased the amount of data that analysts must sort through, and it consequently changed the nature of their jobs.

"Today your first responsibility as an analyst is to keep track of what's happening right now," former CIA analyst Larry Johnson said during a brief phone conversation as he prepared to depart for Iraq on a consulting assignment this past May. "That means dealing with 1500 to 2000 messages, classified at various levels, that move across your desk every day, messages which can be one to three pages long."

Though the volume of the data is greater than it ever has been, the methods for analyzing it haven't changed. Gregory F. Treverton, senior policy analyst at the Rand Corp., Santa Monica, Calif., noted during a recent tour of intelligence agencies that analysts don't use formal analytical methods, let alone computational ones.

"Insofar as there was a method in play, it was limited to brainstorming and then looking for evidence and argument that would either confirm or disprove hypotheses," he says. "Maybe that wasn't such a bad way to do the work during the Cold War, but it seems to many of us that it's not the right way to do analytic work now."

An intelligence analyst's routine these days is more like that of a reporter than that of an academic, according to a 2005 ethnographic study for the CIA's Center for the Study of Intelligence. "Basically, on a day-to-day basis, it's like working at CNN, only we're CNN with secrets," one analyst told the study's author, anthropologist Rob Johnston.

The result has been a major shift in the analyst ranks: 50 percent of U.S. analysts have less than five years' experience, according to some estimates. And yet despite all the turnover, Johnston noted a lingering tendency among analysts to look for information to confirm the prevailing hypothesis in their groups or sections rather than challenge it and risk alienating colleagues and superiors. Indeed, it is considered taboo to change "the corporate product line": if the president or his national security team receives an official opinion from an intelligence agency and that agency later radically revises it, trust, status, and ultimately funding are jeopardized.

Besides looking for patterns in evidence that confirm existing theories, Johnston asserts that analysts often use the wrong rules to make predictions or are too focused on one little piece of the puzzle—say, the influence of foreign fighters in Iraq's Anbar Province. That makes it hard for them to integrate all of the different kinds of information necessary to explore how people might behave in a given situation.

"Becoming an expert requires a significant number of years of viewing the world through the lens of one specific domain," writes Johnston (who did not respond to repeated requests for an inter-

**"I'M NOT SURE THAT MODELING IS GOING TO KEEP THE OPPONENT FROM STRIKING....WHAT WE REALLY NEED TO FIND IS SOME METHOD FOR DETERMINING TERRORISTS"**

view). “This concentration gives the expert the power to recognize patterns, perform tasks, and solve problems, but it also focuses the expert’s attention on one domain to the exclusion of others. It should come as little surprise, then, that an expert would have difficulty identifying and weighing variables in an interdisciplinary task, such as forecasting an adversary’s intentions.”

**SO, CAN COMPUTERS** running agent-based models help analysts fill the prognostication gap and assist them in imagining what they couldn’t otherwise envision?

Barry Silverman, the professor at Penn, thinks so. He suggests that analysts could begin to explore an extensive range of potential terrorist behaviors by having computers model what an individual terrorist wants the world to be and what he is willing to do to achieve that vision.

“Terrorists are not pure evil,” says Silverman, director of Penn’s Ackoff Center for the Advancement of Systems Approach. “They have a value system, and within that value system they are very consistent. If you re-create their value system, their leaders, what they consider to be sacred, what they consider to be violations, if you build up that worldview and that framework, it’s not surprising or crazy that they do what they do.”

Each of the individual agents in a Silverman simulation is an astoundingly sophisticated amalgamation of more than 100 models and theories from anthropology, psychology, and political science, combined with empirical data taken from medical and social science field research, surveys, and experiments [see chart, “Getting Inside the Mind”]. His 15-person research team imbues agents with detailed physiologies that respond to hunger, fatigue, and stress, as well as with complex reasoning skills, long-term memories, and value systems.

Silverman’s students translate theories and models expressed in ordinary language—such as theories of emotions and the origins of stress—into mathematical statements, rules, functions, and algorithms suitable for encoding in the simulations. Using the Python programming language, they have created 500 “performance moderator functions,” each composed of tens to hundreds of lines of code, which represent physical stressors such as ambient temperature, hunger, and drug use; resources such as time, money, and skills; attitudes such as moral outlook, religious feelings, and political affiliations; and personality dispositions such as response to time pressure, workload, and anxiety.

With the agent-modeling editor created by his group as part of its software package, Silverman stitches all of these performance moderator functions together to imitate a consciousness *in silico*. In other words, one of Silverman’s typical agents perceives a stimulus—a gun pointed in the face, a piece of chocolate offered by a soldier—and, through a minutely detailed process that seeks to capture the essence of the workings of the human mind, decides how to react to achieve immediate objectives.

As Silverman explained to a visitor in his office this past May, his software represents each of the various factors that influence perception and reaction as a receptacle, or tank, that can be filled and emptied. So to see how an agent—be it a terrorist, a soldier, or an ordinary citizen—will react when it’s “hungry,” you empty its stomach tank. If you want to give your agent an attribute or stressor, say, alcohol, you just drag a new tank into the agent’s profile and fill it to a certain level to see how different degrees

of a condition, in this case inebriation, might affect emotional and physical responses and, ultimately, decision making.

Such simulated psyches exact a surprisingly modest price in computation and storage. That’s because most of Silverman’s simulations—with the notable exception of the Bakhara market simulation, which runs on a three-dimensional game platform—are practically graphics-free. For instance, on a laptop with a 1.7-gigahertz CPU and 512 megabytes of RAM, Silverman has run simulations involving 1000 agents, each represented by a small triangle that can be observed moving around a town in real time, navigating the streets. On occasion they might flock together to form a spontaneous protest. On a state-of-the-art PC, his group is now simulating a small society of about 15 000 leader and follower agents organized into tribes, which squabble over resources.

To give his agents humanlike behavior, Silverman’s model employs a simple equation. To figure out how an agent will react in a given situation, the software considers a number of possible responses. For each, it takes the expected probability of the response’s being successful and multiplies it by the payoff, or the usefulness of that action in helping the agent achieve its goals. The response that gets the highest score is then chosen. An artificial intelligence expert would use this same equation to create a software robot that can take a set of

facts as inputs—a car is on fire, other agents are fleeing the immediate vicinity, explosions sound in the distance—to deduce that, in this case, there is an attack of some sort going on, and that to achieve its goal of staying alive, its best response is to run.

But for Silverman, the calculation is just the beginning. “AI researchers won’t worry about how stressed the person is, how tired the person is, or what their emotions are,” he says. “Consciousness is a mind-body problem. If I’m really stressed, and I can’t stop and think of

all of the alternatives, I’ll just pick something, and it won’t be the best.”

To endow his agents with more realistic coping mechanisms, Silverman is relying heavily on concepts developed by psychologists Irving Lester Janis and Leon Mann in the 1970s. The Janis-Mann model describes five coping styles that depend on your stress level. When there’s no stress, you’re bored and inattentive. A little stress and you’re on autopilot—it’s Saturday and you’re running through your to-do list. More stress and you reach the perfectly vigilant mode of an expert thinker, where you consider all alternatives before making a decision. Crank up the stress another notch, and you wind up denying the most threatening thing in your existence—think of a terrified teenager in a slasher flick who sees a friend horribly killed and screams “This can’t be happening!” The top level of stress induces full-on panic—you stop thinking and either flee or cower in place.

Emotional responses also figure into human behavior, of course. Silverman found a viable emotion model in *The Cognitive Structure of Emotions* (1988), by Andrew Ortony, Gerald L. Clore, and Allan Collins. According to this model, emotions are feelings about your long-term preferences for the state of the world, short-term goals, and standards of behavior and conduct. But Ortony and his colleagues didn’t discuss how to derive those goals, standards, or preferences. So using data about real people, Silverman and his students crafted a kind of decision tree—a graph of options and

## THE ORGANIZATIONAL RISK ANALYZER USES STATISTICS TO EXAMINE A NETWORK AND DISCOVER MORE ABOUT ITS AGENTS AND HOW THEY INTERACT

associated risks and possible consequences—to model an agent’s near-term goals, standards of behavior, and long-term preferences for how the world should be. Collectively these are called *values*, which, as it turns out, get to the heart of what it is to be a terrorist.

“A jihadist terrorist has a preferred state of the world: the whole world is fundamentalist, and no loose-valued Westerner should walk on the land that Muhammad felt was sacred,” Silverman says. “Those are long-term preferences—control of the land, control of who’s allowed to do what. Then he’s got standards. It’s okay to kill. In fact, I’ll go to heaven and be a martyr, and my family will get money.”

“It’s not game theory,” he adds. “It’s people’s values.”

Silverman doesn’t guess at these values. He works with real-world data provided by leading experts on jihadist terrorists, including Marc Sageman, a colleague of his at Penn. Sageman, a forensic psychologist, wrote a 2004 best seller, *Understanding Terror Networks*, that profiles 172 jihadist suicide bombers. He has since profiled 300 more, including those in the cells responsible for the Madrid train bombings on 11 March 2004.

Sageman’s profile of a typical jihadist terrorist is a male who grows up in a decent family, never gets into much trouble, and goes to a religious school. There he learns Islamic fundamentalist credos, and then in many cases he goes abroad to study—often engineering and often in Europe, where he feels alienated and has trouble finding work. He gravitates to mosques to find comfort in the familiar, and there he meets men with similar feelings of isolation. Gradually, he is drawn toward a politicized existence and into a terrorist cell. By incorporating more than a dozen such models and data sets, some of Silverman’s simulations can emulate how a jihadist agent might feel about his group, when he will become radicalized, and what he might do as a result.


**WHILE SILVERMAN** constructs individual agents from the inside out, sociologist Kathleen Carley peers down on vast groups of agents from on high, building worlds in bottles using publicly available data and classified information. Carley’s models are examples of the network-oriented approach to computer-aided analysis. They help intelligence agencies discover and track a terrorist network’s key operatives, resources, and communication links. Her software can identify leaders of terrorist organizations and let analysts see how a particular network might change if these individuals were arrested or killed—“isolated,” as she puts it. Such dynamic simulations will also tell who might try to take a deposed leader’s place, presumably so that, in the interests of totally dismantling a network, those people could be isolated as well.

BILL GRAMEN/WONDERFUL MACHINE

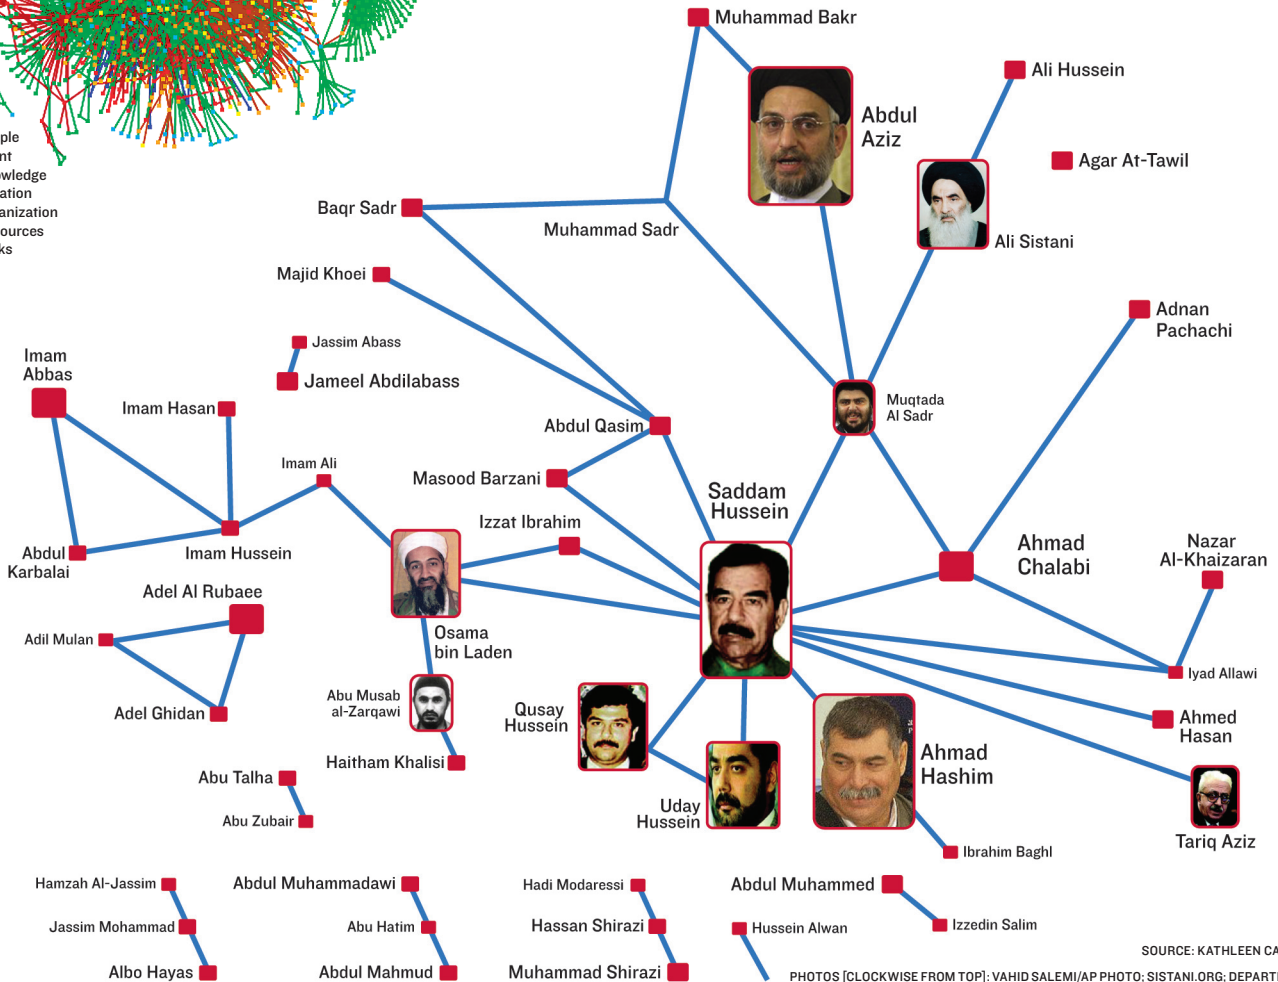


**KATHLEEN CARLEY** designs software programs to help intelligence analysts target terrorists.

**GETTING INSIDE THE NETWORK:** Kathleen Carley and her group at Carnegie Mellon University model a terrorist network in terms of operatives, resources, communications, and so on. [At left is the radical compression of one such model, representing Iraq's Diyala region.] The expansion below shows connections among suspected insurgents, their contacts, and their sources of inspiration, drawn from publicly available data. The size of a box is proportional to the likelihood that the individual is directing others and telling them what to do. The larger the box, the more likely the person is a leader. As this map was constructed just after the U.S. invasion in 2003, Saddam Hussein and his sons Qusay and Uday were still at large and heavily connected, while Abu Musab al-Zarqawi, who would come to lead al-Qaeda in Iraq, was barely a blip on Carley's virtual radar.



- People
- Event
- Knowledge
- Location
- Organization
- Resources
- Tasks



The counterterrorism software package developed by Carley's 30-person team at Carnegie Mellon University, in Pittsburgh, consists of three basic tools: a program that constructs social networks from text files such as newspaper stories and intelligence reports; one that statistically profiles networks in terms of subgroups, individuals, resources, and communications; and a multiagent simulator that shows how social networks evolve over time. Linked together, these programs form a software suite that ingests, analyzes, and models evolving groups and events.

Carley's lab, the Center for Computational Analysis of Social and Organizational Systems, functions more like a small business than a university research group. For example, it has products and works with clients such as the Army Research Laboratory, DARPA, the Office of Naval Research, and the CIA to tailor its software packages to users' specific needs.

With her soft-spoken manner and piercing blue-eyed gaze, Carley projects an authority that must serve her well in the

defense and intelligence communities she inhabits. She holds dual undergraduate degrees in economics and political science from MIT and a Ph.D. in sociology from Harvard. Carley's bookshelves bulge with social science and computer science texts, the manifestation of a career spent trying to blend the two disciplines.

Her expertise and her software are much in demand these days. She hurriedly tells a visiting reporter she can spare only an hour before she'll have to leave on a two-week tour of various three-lettered government agencies.

She pulls out a laptop to show how one of her group's programs, AutoMap, ingests text from Web pages, newspapers, magazines, intelligence reports, and interview transcripts. The program examines chunks of text for names and words contained in its user-defined thesaurus that are within a certain proximity to one another—for example, within the same sentence, paragraph, or article. By linking people, resources, and events, AutoMap creates multilayered networks organized according to relationships among people, resources, knowledge, locations, and events.

For instance, from a single newspaper article, AutoMap detailed the relationships among eight people linked to the terrorist group Jemaah Islamiyah (JI) and the 2002 bombings in Bali that killed 202 people. Using the relative proximity of terms in sentences, AutoMap created links between individuals. Take this sentence: “They would later find during the interrogation of Mr. Rusdi that walking several meters behind Mr. Rusdi when he was arrested was Azahari Husin, a Malaysian lecturer and one of two key JI bomb-makers who remain at large.”

From the proximity of the names Rusdi and Azahari Husin, AutoMap infers that a relationship exists between the two men and between them and Jemaah Islamiyah. It also plucks out relevant phrases such as “Malaysian lecturer” and “JI bomb-makers” and relates them to these two people, and by association to the people to whom they are connected, to create a map of the JI network.

Like Silverman’s team, Carley’s group has used Iraq as an active test bed for its programs. On the laptop screen, Carley pulls up a social network constructed by AutoMap from publicly available news sources of the Diyala region, just northeast of Baghdad [see chart, “Getting Inside the Network”].

Carley points to the screen, which displays a virtual cobweb of dots linked by lines. “Here’s the resources network and how they’re connected,” she says. “These are the organizations and how they’re connected, financial as well as weapons,” and many other resources like food, clothing, lumber, and fuel.

Next, she runs a subprogram within AutoMap called the intel report, which in the case of the Diyala region that she’s got on her monitor identifies the emerging leaders of the Iraqi insurgency according to rank, who’s interacting with whom, who’s got what knowledge and resources at their disposal, and who’s been seen in what locations.

AutoMap spits out the network in a form of eXtensible Markup Language (XML) and passes the data on to another program written by Carley’s group, the Organizational Risk Analyzer. It uses statistics to examine a network to discover more about particular agents and how they interact with each other and influence group dynamics. Based on network theory, social psychology, operations research, and management theory, the software identifies the most important agents in a terrorist network.

Before identifying key individuals, however, first you have to find the covert network, which is by definition working hard to stay hidden. Here the Fuzzy Overlapping Grouping (FOG) algorithm, developed by Carley’s grad student George B. Davis, provides something like an X-ray view of a society.

A number of programs can uncover networks of people buried in reams of telephone records, e-mails, and Internet traffic logs. But all of them assume that each person is a member of a single group that has no overlap with any other groups and that all people in a group are equally dedicated to its cause. Yet, as people have observed in Iraq and elsewhere, alliances of individual terrorists constantly shift among groups. For instance, an aspiring terrorist might be plotting to kidnap a foreign journalist, but the group he initially belonged to has begged off, so he joins one that will help him carry out the operation. As Davis explains, the FOG algorithm is unique in that it takes such morphing and multiple allegiances into account.

On the screen, FOG’s results are inscribed as lines connecting individuals into groups, lines whose thickness is determined

by the degree of an individual’s dedication to a particular group. Individuals with multiple ties are called *interstitial people*, and they are often group leaders. Once these individuals are identified, analysts devise hypothetical interventions, perhaps an arrest or air strike, to isolate those individuals from their groups.

To see how such interventions might play out, the risk analyzer exports its statistical analysis of the network and its agents, up to 10 million of them, to the simulation program called DyNet. The user then defines the network’s mission and technological capabilities for DyNet: it might postulate that a group wants to blow up an oil pipeline, uses cellphones to communicate, but doesn’t have the skills to make a remote-controlled bomb. Next, the user chooses a scenario for DyNet to simulate: what if the person who is responsible for relaying messages from the group leader to the group members is captured? DyNet will then simulate how the leader and his followers would try to reconnect.

**BECAUSE LARGE MODELS** of an entire city or region can take hours to run, even on a supercomputer, Carley demonstrates DyNet on her laptop using a much smaller sample: three research groups in her lab, a mere 31 agents, all connected through one interstitial agent, her.

According to Carley, people connect for two basic reasons: because they are similar to each other and because they need information. “If something breaks the network and an agent realizes it’s cut off, it’ll start trying to interact with others on the basis of those two rationales,” she says. “And eventually someone will respond.”

She eliminates herself on screen to show how individuals in the groups now cut off from each other start sending out feelers to find the other group. The groups eventually reconnect according to the social theories and data coded into DyNet.

Carley, who knows Silverman’s work well, characterizes their approaches in yin-yang terms. “We do a lot more with networks; he doesn’t do networks as much,” she explains. “Our models are more macro social networks; his are more of a single person with incredibly complex attributes and theories of mind. His agents have a much more emotional affective component. We hardly ever put emotions into our models,” she added.

She hesitates for a moment before saying, “The real breakthrough in the next 15 years is when we marry Barry’s work and mine.”

It’s a vision that carries with it both tremendous potential and possible risk. A virtual world in which autonomous agents simulate precisely how individuals and networks might react to even subtle changes in social, political, and economic conditions could have many uses besides figuring out what terrorists might do next.

On the bright side, computers could do the guesswork so people can execute complex projects, involving multiple networked groups, which will have a better chance of succeeding than they do now. When planning to radically restructure a national government around, say, the mission of securing the homeland, policy wonks could use multiagent models to tease apart the complex interdependencies among disaster response planning, intelligence gathering and sharing, law enforcement, the federal bureaucracy, and the private sector. Using actual data, they could model how people might react to changing job functions and discover the

**“A SIMULATION IS BY ITS NATURE SPECULATIVE, AND YOU DON’T GO OUT AND KILL PEOPLE BASED ON SPECULATION”**



optimal way to reorganize a complicated bureaucracy so that it executes its core mission as effectively as possible.

Then there's the dark side: imagine a Bland New World where television shows, books, movies, music, fashion, and food are carefully calibrated to the simulated desires of a mass audience, shutting out marginalized artists and depriving consumers with marginalized tastes of easy access to the pleasures they crave. And then there is the precognitive dystopia famously imagined by the science fiction writer Philip K. Dick in his 1956 short story "The Minority Report"—just replace Dick's three gibbering psychics who predict future crimes with data centers full of ultrapowerful computers that do something similar, only without the gibbering.

**OF COURSE, THESE SCENARIOS** might never play out in reality, particularly if today's models don't spread beyond the handful of techno-savvy analysts who have the inclination and the patience to tune them for their particular needs.

"It would be nice to do these simulations, and they're probably fun to do, but unless you can demonstrate an immediate benefit to justify the cost in terms of an analyst's time, it's a luxury the analyst can't afford," former CIA analyst Larry Johnson says. "The bottom line is that these tools probably won't be used."

Analysts feel that they can't fit computational modeling into their routines, because many of these programs aren't designed with the end user's needs in mind, argues Jim Nyce, an anthropologist at Ball State University, in Muncie, Ind.

"This is a particularly homegrown American phenomenon, because of the strength of the modeler community here," Nyce says. "But there is at least anecdotal evidence that in our civilian intelligence community in the States a lot of the good, scientifically robust modeling environments don't get used because the cognitive, intellectual, and work requirements have not been taken into account in their design."

Nyce, who is also a visiting professor at the Swedish National Defense College, in Stockholm, points out that Europeans, and Swedes in particular, design desktop tools by first looking at the kinds of tasks that someone in the military or emergency response community would have to handle. "And in the U.S., the emphasis seems to be if we produce a strong modeling environment, the problem is solved. People will just naturally use it," he observes. "But when you're working with high-volume, data-intensive job tasks, you can't do a lot of on-the-job training. So no matter how powerful a resource may be, if it doesn't fit with your work requirements and information requirements, the chances are no one's going to click on it on their desktop."

Carley, who has designed her suite of programs specifically for the desktop, insists that she has done so in partnership with working analysts. But she concedes that right now her software requires huge amounts of effort and time to input the rules and data that drive the simulations and still more time to learn to use.

"We're trying to address those kinds of issues," Carley says of both her group and modelers in general. Her long-term goal is a system so straightforward that analysts can, "with a week's worth of training, become well enough versed to understand what they're really seeing."

But how certain can analysts be that what they see on screen accurately reflects reality? Carl Mitcham, a philosopher of technology and editor of the *Encyclopedia of Science, Technology and Ethics*

(Macmillan, 2005), contends that no matter how advanced the model, it will leave out nuances that tell us why people act the way they do. Models, he argues, can't possibly get at the deep mystery of why someone straps on a vest of explosives and wanders into a crowded restaurant or flies a jet plane into a building filled with people he doesn't know.

"I fully recognize that human behavior modeling has become increasingly sophisticated," says Mitcham. "But such approaches don't necessarily contribute to my understanding the reasons that might be motivating the behavior or how to transform the individual whose behavior I'm modeling."

From Mitcham's point of view, modeling allows, and maybe even encourages, people to avoid considering the legitimacy of certain grievances that drive terrorist behavior. Without that understanding, he argues, the West won't be able to effectively win over the hearts and minds of adversaries, and as a result, the world will be locked into a perpetual cycle of attacks and counterattacks.

"I'm not sure that modeling is going to keep the opponent from striking. It will only help you in killing some of the opponents," Mitcham concludes. "What we really need to find is some method for deterring terrorists rather than just killing them."

**"A JIHADIST  
TERRORIST HAS  
A PREFERRED  
STATE OF THE  
WORLD.... HE'S GOT  
STANDARDS....  
IT'S NOT GAME  
THEORY. IT'S  
PEOPLE'S VALUES"**

Whether the goal is to kill, deter, or negotiate with terrorists inside a computer or in the real world, ultimately agent-based models must answer one fundamental question: Do they help save lives?

As with so many intelligence projects whose existence never leaks beyond the inner sanctum, the general public will probably never learn whether a certain agent-based model was used to disrupt a specific terrorist plot. "It wouldn't be known, frankly," says Charles Macal, director of the Center for Complex

Adaptive Agent Systems Simulation, at Argonne National Laboratory, in Illinois, whose NetBreaker social network model is currently being tested by analysts. "The CIA or the [Defense Intelligence Agency] isn't going to publicize anything."

A chance to test this hypothesis came on 7 June of this year, when two 500-pound bombs obliterated a house near Baquba, in the Diyala region of Iraq, killing six people, including Abu Musab al-Zarqawi, the leader of al-Qaeda in Mesopotamia. U.S. and Iraqi forces immediately began conducting raids on people with suspected ties to Zarqawi's network. One of Carley's models is of the Diyala region. Was there a link to the raids?

Carley's e-mail reply to this query came almost three days later. "Sorry," she wrote. "No comment." ■

## TO PROBE FURTHER

Rob Johnston's ethnographic study "Analytic Culture in the U.S. Intelligence Community" details the workaday world of today's intelligence analyst: <http://www.fas.org/irp/cia/product/analytic.pdf>.

The melding of computer modeling and social science is just getting started. To hear presentations from the first two National Security Threats conferences held at the Joint Threat Anticipation Center at the University of Chicago, go to <http://jtac.uchicago.edu/conferences/05>.

To see video clips of Barry Silverman's *Black Hawk Down* simulation, check out <http://www.seas.upenn.edu/%7Ebarrystg/HBMR.html>.

More about Kathleen Carley's research is available at <http://www.casos.cs.cmu.edu>.