Out of the Ordinary

Finding Hidden Threats by Analyzing Unusual Behavior

JOHN HOLLYWOOD, DIANE SNYDER,
KENNETH McKAY, JOHN BOON

Approved for public release, distribution unlimited

RAND
CORPORATION

This research in the public interest was supported by RAND, using discretionary funds made possible by the generosity of RAND's donors, the fees earned on client-funded research, and independent research and development (IR&D) funds provided by the Department of Defense.

Library of Congress Cataloging-in-Publication Data

Out of the ordinary : finding hidden threats by analyzing unusual behavior /
John Hollywood ... [et al.].
p. cm.
"MG-126."
Includes bibliographical references.
ISBN 0-8330-3520-7 (pbk. : alk. paper)
1. Criminal behavior, Prediction of—United States. 2. Crime forecasting—United States. 3. Criminal methods—United States. 4. Terrorism—Forecasting. 5. Terrorism—Psychological aspects. 6. Intelligence service—United States. 7. National security—United States. I. Hollywood, John S., 1973– II. Rand Corporation.

HV6080.O97 2004
363.32—dc22

2003023703

Cover photograph by Kenneth N. McKay. The photograph is of the "Warabe-Jizo" statue in the Yusei-in Garden of the Sanzen-in Temple in Ohara, Japan. The statue is of a child bodhisattva-kshitigarbha. He is a figure from both the Hindu and Buddhist religions. Derived from the Mother Earth, he appeared in the world to help people.

The RAND Corporation is a nonprofit research organization providing objective analysis and effective solutions that address the challenges facing the public and private sectors around the world. RAND's publications do not necessarily reflect the opinions of its research clients and sponsors.

Published 2004 by the RAND Corporation
1700 Main Street, P.O. Box 2138, Santa Monica, CA 90407-2138
1200 South Hayes Street, Arlington, VA 22202-5050
201 North Craig Street, Suite 202, Pittsburgh, PA 15213-1516
RAND URL: http://www.rand.org/
To order RAND documents or to obtain additional information, contact
Distribution Services: Telephone: (310) 451-7002;
Fax: (310) 451-6915; Email: order@rand.org

Summary

The problem of "connecting the dots" in intelligence—selecting and assembling disparate pieces of information to produce a general understanding of a threat—has been given great priority since the September 11, 2001, terrorist attacks.[1] This monograph summarizes a RAND internal research and development project on developing unique approaches to assist in connecting the dots.

Synthesizing disparate pieces of information to understand threats is an extremely difficult challenge. The analysis process requires searching through enormous volumes of data, and analysts' attention must be directed to the most important findings. There are, however, few direct clues as to which data are important and how to link the data together. The most obvious approach to prioritizing data—looking for patterns similar to those of previous attacks—can easily lead to missing the signals indicating the next, different attack. When analyzing uncertain and messy (i.e., real-world) data, time and situational pressures often force the analyst into making conclusions, despite great uncertainty as to whether the conclusions are true. Ex-

[1] As one example of the high priority placed on this topic, the Congressional Joint Inquiry into September 11 writes, in its "Conclusion—Factual Findings" section: "No one will ever know what might have happened had more connections been drawn between these disparate pieces of information. We will never definitively know to what extent the Community would have been able and willing to exploit fully all the opportunities that may have emerged. The important point is that the Intelligence Community, for a variety of reasons, did not bring together and fully appreciate a range of information that could have greatly enhanced its chances of uncovering and preventing Usama Bin Laden's plan to attack these United States on September 11th, 2001."

isting legal, technological, procedural, and cultural barriers to sharing and linking information further complicate these challenges.

A Schema for Connecting the Dots

Historically, however, many people have surmounted the barriers to connecting the dots, albeit with significantly smaller amounts of data than the homeland security community faces. These successful problem solvers have tended to follow similar cognitive processes. First, the problem solver establishes expectations for what the environment will be like if everything is "normal"—in effect, defining a *status quo*. This formulation is employed because it is often impossible to predict everything that is abnormal; instead, it is much easier to describe the status quo as the starting point and add to this description what is known about how the status quo might change. The problem solver next identifies a set of metrics (both quantitative and qualitative) with which to observe the environment, especially in regard to whether the actual environment is consistent with expectations. Third, the problem solver observes streams of measurement data about the environment. Generally, the solver does not examine every observation carefully but instead scans for *out-of-the-ordinary* or *atypical* signals that significantly deviate from the expected status quo. These signals range from defined precursors of a well-understood change in the environment to an entirely novel phenomenon whose meaning is unknown—except that it is in some way relevant to the task at hand.[2] All, however, deserve additional analysis: Because they are outside of expectations for what the current environment should exhibit, they

[2] It is important to reiterate that the problem solver does not try to examine all atypical behavior in the environment; doing so would lead to data overload. Instead, the solver pays attention to relevant behavior that can quickly be related to the task at hand. For example, suppose the problem solver is responsible for identifying potential threats to a theme park. Clearly, many attendees in the theme park will engage in "unusual" behavior. The problem solver, however, will be interested strictly in behavior that can quickly be declared potentially relevant to attacks on the theme park, such as a group of guests on a terror watch list, or a group of guests who engage in behavior that strikes the park's security guards as threatening (casing behavior, clandestine communications, etc.).

may signal an impending change in the environment. Upon discovering out-of-the-ordinary behavior, the solver looks for supporting data marking the observed signals as a true phenomenon and not just noise. Should such supporting data be discovered, the problem solver searches for related information that helps explain the phenomenon and then develops and tests hypotheses as to what the phenomenon means. Finally, once the phenomenon is understood, and identified as indicating a risk, the problem solver uses heuristics to avoid or mitigate the risk. It should be noted that the process the problem solver uses is not linear—the solver separates the noise from the truly significant through an iterative, multistage process of testing and learning, with the steps used being dependent on what the solver learns about the phenomenon at each stage (i.e., *context-dependent* analysis).

We have developed the **Atypical Signal Analysis and Processing** (ASAP) schema to assist in connecting the dots by mirroring the problem-solving process described above. An implementation of the schema will serve as an analyst's "virtual extension," applying the problem-solving process to the volumes of data and numbers of dimensions within the data that are far too large for analysts to work with directly. Figure S.1 shows the schema.

The shortest, linear path through the schema has six major steps. The schema begins with the gathering of information from a set of external databases. Most of the information pertains to *watched entities*—people, places, things, and financial activities already suspected as being relevant to a terror attack or activities within key infrastructure and commercial processes already being monitored, such as international commerce, nuclear energy, hazardous materials, and air transportation. Intelligence and government databases would be used, supplemented by open-source data, all in accordance with privacy regulations. This baseline information would be further supplemented by *precedent-setting phenomena*—data, voluntarily submitted, that describes behavior the reporters find to be highly out of the ordinary and suspicious with respect to asymmetric threats. (For ex-

Figure S.1
The Atypical Signal Analysis and Processing (ASAP) Schema

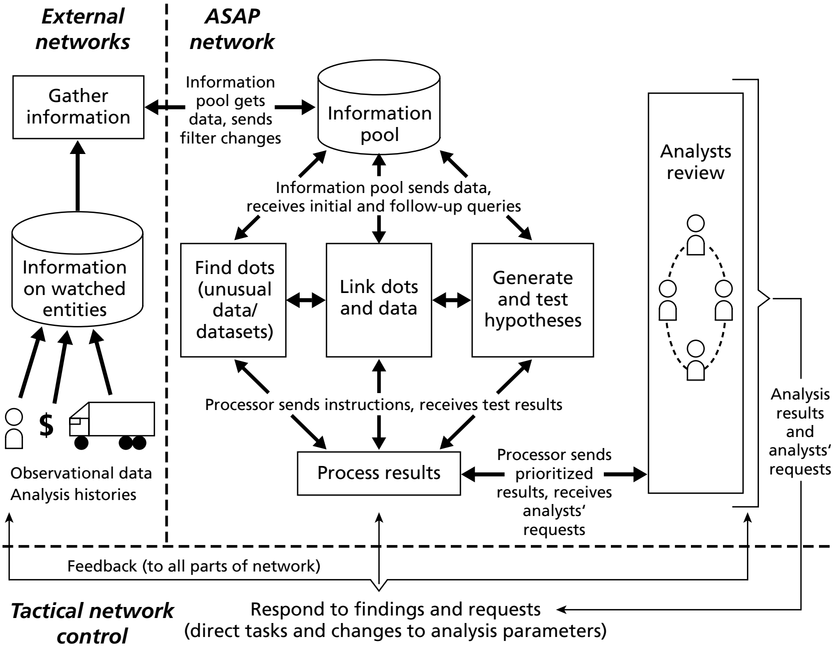
External networks
ASAP network
Gather information
Information pool gets data, sends filter changes
Information pool
Information on watched entities
Information pool sends data, receives initial and follow-up queries
Find dots (unusual data/ datasets)
Link dots and data
Generate and test hypotheses
Analysts review
Processor sends instructions, receives test results
$
Observational data
Analysis histories
Process results
Processor sends prioritized results, receives analysts' requests
Analysis results and analysts' requests
Feedback (to all parts of network)
Tactical network control
Respond to findings and requests (direct tasks and changes to analysis parameters)

RAND*MG126-S.1*

ample, prior to the 9/11 attacks, FBI officials might have submitted their suspicions about certain flight school students.) The schema incorporates both direct observations of the watched entities and metadata on who is working with those observations and why. The resulting information goes into a structured information pool.

Second, within the pool, a number of automated detection agents perpetually filter the information to look for out-of-the-ordinary signals.[3] These signals might be single observations (e.g., a

[3] Note that an ASAP network would not detect and process all atypical signals; instead, it would process atypical signals that can be quickly classified as being potentially relevant to an attack or the operations of a terrorist organization. For the former, a network would seek atypical signals potentially related to attack preparations such as target casing, training, clandestine communications, supply (smuggling), and weapons acquisition. For example, from a theme park, the network would be interested in hearing reports of people videotaping secu-

very large financial transfer) or a significant trend (e.g., a 75 percent increase in fund transfers during the past month). The signals might also be a group studying information they do not normally review (e.g., an FBI field office requesting records of students at truck driving schools funded by the aforementioned increase in funding transfers). Such signals become the "dots." Note that ASAP will support detection filters ranging in sophistication from simple rules evaluating a few data fields (usually generated by human analysts) to complicated algorithms evaluating tens of simultaneous data fields simultaneously (usually generated by hybrid human-machine statistical training techniques, such as neural networks).

Third, once the dots have been identified, the next step is to find information related to the dots. The schema thus employs automated relationship agents to look for relationships between new and existing dots. It also uses agents to perform *backsweeping*—searching for previously unremarkable data that relate to the dots. These related data would come primarily from the information pool but also from queries in external (intelligence) databases and, in cases constituting probable cause, from commercial databases (for example, examining the credit transactions of a positively identified terror suspect).[4] The information discovered helps determine the extent of an out-of-the-ordinary phenomenon and provides a context to help explain it.

Fourth, once the dots have been linked, hypothesis agents can be tasked to create possible interpretations for the linked dots and to create corresponding testing plans to determine whether the hypotheses are correct. The principal purpose of these agents is to assess which phenomena should be given priority for further investigation.

rity checkpoints and support beams of major attractions; it would not be interested in hearing reports on generic disorderly conduct. For the latter, a network would seek atypical signals such as sudden movements, changes in organizational structure, or changes in communications networks. The issue of what constitutes "out of the ordinary" is discussed at length in Chapter Two.

[4] Backsweeping in probable-cause cases is the only time the ASAP schema would use general commercial databases. Thus, for example, the schema complies with the proposed Citizens' Protection in Federal Databases Act, which would prohibit accessing databases "based solely on a hypothetical scenario or hypothetical supposition of who may commit a crime or pose a threat to national security."

Consequently, the "hypotheses" very often do not pertain to a specific inference but instead simply note that a phenomenon is so unusual (and perhaps has particularly suspicious characteristics) that it is worth investigating further. Correspondingly, the testing agents monitor whether further investigations raise or lower concern about the phenomenon.

Fifth, the results of these processes are strictly prioritized, and high-priority results are forwarded to analysts. This prioritization function is one of the most important of the schema, as it reduces potentially large volumes of out-of-the ordinary discoveries, so that analysts can restrict their attention to only the most relevant and significant discoveries.

Finally, the schema facilitates the collaboration of analysts working on related observations. It notifies different analysts that they are looking at the same pieces of information and provides communications channels between them. In the ASAP schema, analysts have primary responsibility for actions to be taken in response to unusual phenomena that are brought to their attention because they have insights (knowledge of human behavior, for instance) that automated systems do not have.

As with human problem solvers, the schema permits iterative, dynamically tailored analysis in which the actual sequences of testing activities are dependent on what has been learned to date about the observed phenomena. To allow for such context-dependent processing, the complete schema is governed by a two-stage control system. At the lower, operational level, processor agents direct data through the schema. These agents use sets of control rules to interpret the results from the detection, relationship, and hypothesis agents, and determine what to do next with a particular dataset (or test results on the dataset). Thus, for example, a processor agent might direct a newly detected dot to a relationship agent and forward results from hypothesis testing to analysts. This structure allows for flows through ASAP to be both dynamic and iterative. Thus, analysis results guide what happens next, so that, for example, analyzing one initial signal leads to the discovery of related phenomena, which are then further analyzed, leading to yet more contextual information, and so on, potentially allowing an initially mysterious phenomenon to be illuminated fully. Processor agents are guided both by automated logic and directions from analysts. Analysts have the ability to request any type of follow-up test or analysis of the ASAP agents, with the processor agents executing these requests.

At the second, tactical level, the ASAP is subject to open-loop control: Analysts may change any of the software agents and agents' parameters, or make any specific analysis requests, in response to the analysis results. The tactical level also supports automated control agents that modify software agents and parameters based on interpretation of finding, relating, and testing dots (these software control agents are also subject to analysts' direction).

We have developed an architectural design that applies the schema; description of the design makes up the bulk of this paper. The design has several key attributes worth mentioning here.

First, in its initial stages the architecture focuses on information already suspected of being of interest, as opposed to performing unguided data mining of large databases and collecting data about generic transactions. This focus helps prevent analytic overload. At the same time, the architecture has the flexibility both to receive reports of highly atypical behavior from all sources and to cull databases for particular pieces of information should the need arise (for example, searching for data about a highly suspicious person's travel plans).

Second, the architecture searches primarily for signals that are out of the ordinary as opposed to signals that fit predetermined patterns. This approach loses precision in meaning but gains in being able to detect a wide range of threatening behavior that does not fit previously seen attack patterns. Searching for signals deviating from, rather than matching, existing patterns is uncommon in the pattern-matching and signal analysis fields.

Third, in finding dots, searching for related information, and generating hypotheses, the architecture employs contextual rules that allow data to be analyzed in the context of existing knowledge. Contextual rules are not commonly used in information analysis.

Fourth, the architecture explicitly deals with uncertainty by generating and testing competing hypotheses for unusual signals. This approach helps defend against prematurely accepting an explanation for a phenomenon.

Finally, the architecture enables the collaboration of personnel needed to connect the dots, even if the personnel are distributed across different groups and agencies. The architecture looks not just for out-of-the-ordinary data, but for *out-of-the-ordinary analyses of the data*. Flagging these analyses can bring together groups of people and automated agents who can jointly characterize a previously mysterious phenomenon.

Near-Term Implementation

Fully implementing the ASAP schema and its supporting architecture would be a lengthy, multiyear process. However, several improvements could be implemented quickly, in effect allowing personal analysis interactions to partially substitute for the automated agents described previously.

A major requirement for detecting out-of-the-ordinary phenomena is to understand what constitutes "ordinary" and what types of behaviors are significant deviations away from the ordinary that may be relevant to a counterterrorism investigation. Thus, we recommend that appropriate users throughout the homeland security (HLS) community create and distribute standardized profiles of organized behavior. These profiles would discuss both what threats (terror attacks, terror support activities, etc.) commonly look like and what status-quo conditions look like in such "watched" fields as international commerce, transportation, and demolition. Note that these brief profiles are in no way intended to be comprehensive; their purpose is merely to help analysts and field professionals in one area educate analysts and field professionals in other areas—in a more intentional and systematic way than at present—on what types of behavior to look out for.

The next step would be to establish electronic posting boards where those in the field can report unusual phenomena and see whether others have been observing similar or related occurrences—in effect, helping each other serve as detection and linking agents. Personnel would post to unmoderated electronic bulletin boards, and there would be no approval process for phenomena posted. Trained reviewers would routinely review the boards, selecting especially unusual and significant reports to post to filtered boards that would be widely read by analysts.

The third step would be to develop semiautomated tools to help HLS personnel identify posts relevant to what they have been observing. One might first implement organizational tools that divide the posts into threads dedicated to particular occurrences and create indices of those threads. Particularly important threads would be associated with journals or diaries summarizing key developments and current hypotheses. The next step would to be create Google-like search engines for posts that match the results of search queries. Finally, simple heuristics could be developed that look for connections and patterns across the threads of posted messages.

Summarizing the Schema

Table S.1 summarizes differences between the proposed schema and traditional methods of intelligence analysis. The table also compares a near-term, manual implementation of ASAP with a full implementation.

A Research Plan

At the same time as the short-term improvements are being implemented, research can begin on the automated portions of the ASAP architecture. This portion will be needed to assist analysts in identifying out-of-the-ordinary signals in the enormous volume of data generated by intelligence and infrastructure collection and monitoring systems every day.

Table S.1
The ASAP Schema

Traditional Analysis	ASAP Advantages	ASAP Near-Term Implementation	Full ASAP System Implementation
Focuses on previous patterns	Searches for out-of-the-ordinary behavior, allowing for detection of previously unseen threats	Core or pilot group	New communities added to electronic boards
Time pressure drives toward premature closure	Allows memory of hypotheses and data rejected by analysts	Drafting short profiles of existing asymmetric threats—e.g., suicide bombing	Incorporates entire homeland security community
Analysts mostly operate on basis of own experience and biases	Leaves key analytic choices with analysts	Drafting short profiles of status quo in such watched domains as international commerce	Detailed architecture for out-of-the-ordinary analysis
Search tools mostly weed out what doesn't fit pattern	Notices what analysts are watching and asking	Users post on unmoderated electronic boards	Formal specifications for detection, linking, and hypothesis agents
Analysts are isolated within own groups and agencies	Facilitates collaboration between analysts studying the same phenomena	Moderators connect across analysts and, when possible, organizations	Analysis processes integrated across organizations

The first stage of research should develop a detailed architectural plan for the ASAP system and its constituent control and analysis agents. The architecture would specifically describe detection, linking, and hypothesis agents in such key areas as direct threat detection, international shipping, and air transportation. The first stage should also describe how the architecture would address a detailed terror-attack scenario.

The second stage of research should create formal design specifications for the agents and the software making up the ASAP backbone. These specifications would define the objects, methods, and major algorithms employed by the agents and systems management software.

The third stage of research should create a prototype system that would include simple examples of the above agents. It would also include the control components needed to achieve dynamic, feedback-based control. Once the prototype is completed and evaluated, construction and implementation of a real-world ASAP system could commence, moving the ASAP concept from research to reality.