Cognitive Fuzzy Modeling for Enhanced Risk Assessment in a Health Care Institution

Elmé Smith and Jan Eloff, Rand Afrikaans University

HUMAN OBSERVATION AND INTUITION FORM THE BASIS OF ANY RISK ASSESSMENT. THE AUTHORS PROPOSE A COGNITIVE FUZZY-MODELING APPROACH TO RISK ASSESSMENT IN HEALTH CARE INSTITUTIONS, BUILDING ON FUZZY LOGIC'S GREAT POTENTIAL IN DEALING WITH VAGUE INFORMATION AND HUMAN COMMON SENSE AND INTUITION.

ALTHOUGH HEALTH CARE HAS long been lagging behind other industries in adopting computer technologies, the health care industry is now using information technology in almost every health care sector across the globe. This has resulted in major transformations of the entire industry, aimed at maximizing the quality of medical care while minimizing costs.[1] However, storing health care information electronically does raise concerns about the risk of exposing highly confidential and sensitive health care information to outsiders. The health care domain must properly assess the possible risks of computerization and make recommendations for, at best, preventing such risks and, at worst, minimizing them. Because health care information systems are unique, they require a different approach to risk assessment.

Assessing risks

One salient feature that sets health care institutions apart from ordinary institutions is that they are principally aimed at treating people. To incur a risk here (such as unauthorized access to patient information) could compromise not only the patient's privacy but also his or her well-being. Putting patient records online increases the risk that outsiders will exploit sensitive patient data. In the health care environment, other risks, such as the unavailability of patient information owing to a power failure, could have fatal consequences. Also, it is often difficult, if not impossible, to isolate the health care system's assets from the traffic flow of patients, visitors, and doctors, which creates an urgent need to protect patients' privacy.

Just as important is the need to share accurate patient information and ensure its availability to all authorized parties to allow proper treatment of the patient. The dilemma of obtaining, using, and sharing patient information to provide care while not breaching patient privacy is a serious concern.

It is also difficult to quantify nonmonetary risks in a health care environment. For example, it is difficult to determine the cost associated with the incorrect treatment of a patient resulting from inaccurate information. Furthermore, some patient information, such as clinical information, might be confidential, whereas another part, such as geographical information, might be unclassified. The latter, therefore, introduces a certain degree of vagueness regarding the patient

information and the possible risks that a typical health care institution might incur.

We researched several alternative modeling techniques, such as the probabilistic theory, PERT (Program Evaluation and Review Technique) analysis, heuristic modeling, and fuzzy logic. We concluded that fuzzy-logic techniques present a plausible way of modeling such vagueness, and we can relate everything back to a certain degree of likelihood—for example, the chance that patient information will be exposed to outsiders might fall under the "high likelihood" fuzzy set, at a 68% chance.

In addition to vagueness, we must accommodate intuition in modeling risk assessment in a health care institution, because human observation forms the basis of any risk assessment. For example, we cannot precisely determine the likelihood of exposing database files with patient information to outsiders, but we can estimate a value based on observations. Fuzzy logic ensures that we do not neglect human common sense and intuition.

This fuzzy-logic approach lets us calculate IT risk values for different areas in an health care institution. For example, a hospital superintendent can use these IT risk values to identify areas in a specific division that are critical to information security and then enhance the hospital's IT security. Successfully executing this approach, however, requires knowledge-engineering and business-analysis skills. The knowledge engineer needs to acquire the necessary information to, for example, determine whether distributing patient information impacts the patient's privacy while he or she is in the operation ward. Business analysts, on the other hand, should be well educated in the health care profession. They must be able to point out the different sections in a hospital and explain the sharing of patient information during the patient's stay. Medical doctors and professional nurses should be able to assist the business analyst. The knowledge engineer, business analyst, hospital superintendent, medical staff representatives, and administrative staff representatives can all be involved in the knowledge-acquisition process in a typical hospital.

Dynamic health care institutions

When patients visit a typical health care institution, they can follow various routes, depending on the purpose of their visit. Figure 1 depicts a typical example of a patient's route.

Each patient route consists of a finite number of phases. A *phase* is a division within a specified patient route. The example in Figure 1 consists of six phases: registration, preparation ward, operation theater, ward after operation, release, and follow-up visits. To elucidate the assessment of risks through a cognitive fuzzy-logic approach, we use the patient route in Figure 1 as an example throughout this article.

A typical patient spends a certain amount of time in each phase of the route. During each phase, the patient information is essentially shared by authorized communicating parties, such as the doctor, specialist, and laboratory. However, it is also possible that unauthorized parties might need to access the patient information, such as in an emergency situation. If a patient were, for example, admitted to a hospital's emergency room, the on-duty doctor (who is not necessarily a resident at the hospital the patient normally visits) would need to access the patient information at once to effectively treat the patient. The inaccessibility of patient information in such case might have serious consequences. The flow of health care data is, therefore, complex and not limited to the point of care. For this reason, the patient information would also be exposed to many outsiders, thus increasing the possibility of compromising confidentiality, integrity, or availability of the sensitive patient information. In the patient route in Figure 1, we can use different technologies in each phase of a patient's stay, including database and paper files in registration and the preparation ward; microfilm, database files, a LAN server, and paper files in the post-operation ward; and microfilm for the release phase.

TO OBTAIN A CLEAR PICTURE OF THE RELATIONSHIPS BETWEEN THE VARIOUS ASPECTS IN A TYPICAL HEALTH CARE INSTITUTION, WE CAN USE A FUZZY COGNITIVE MAP.

Because these technologies are extremely vulnerable to risks, the institution's staff must consider them when putting in security controls. The security controls will also differ from phase to phase, depending on the possible risks. It might, for example, be easier to gain unauthorized access to paper files than to database files. Therefore, we must be able to revise and enhance security controls continuously for this dynamic environment.

To obtain a clear picture of the relationships between the various aspects in a typical health care institution (such as patient route consisting of phases, time spent in each phase, authorized and unauthorized communicating parties sharing patient information, technologies for storing and processing patient information, and security controls), we can use a fuzzy cognitive map. FCMs are fuzzy-graph structures that provide an expressive and flexible method of intuitively capturing and representing complex relationships.[2] In the event of an intuitive activity such as a risk assessment, the FCM naturally represents the human way of thinking.[3]

Drawing a FCM requires both a knowledge engineer and a business analyst's expertise. They aim to provide a clear representation of the underlying physical, real-world domain with concept nodes representing events that they link to one another. However, the knowledge engineering effort is complex. In the health care domain, the knowledge engineer should be skilled in the interview techniques necessary to acquire raw knowledge from the health care personnel. Furthermore, the knowledge engineer and business analyst need to work closely together to determine how the various dynamic aspects in a specific health care institution influence one another. Finally, they must cast this raw knowledge in an appropriate form that we can use to draw the FCM.

Figure 2 depicts an FCM for a patient's registration phase. The FCM consists of nodes (or concepts) that represent events (for example, the patient spends some time in registration) and edges, which describe relationships (or causal flow) between these events.[2,4,5] These relationships indicate whether one event increases or decreases the likelihood of another event. The edges have "fuzzy" strengths in the interval range [–1,1], indicating the degree to which one event affects another. The plus relationship in Figure 2 between C_1 (communicating parties sharing patient information) and C_6 (the risk of exploited patient information) implies, for example, that if the number of communicating parties sharing patient information in the registration phase increases, then the risk of patient information being exploited in this phase will also increase by a degree of 0.8, or 80%. By the same token, if the number

of communicating parties decreases, then the risk of patient information being exploited will also decrease by 80%. The other plus relationships work in the same way.

On the other hand, the minus relationships indicate that the possibility of one event occurring increases while the possibility of another event occurring decreases, and vice versa. In this way, the minus relationship between C_5 (security controls) and C_6 (the risk of exploited patient information) implies that if the strength of security controls in the registration phase increases, then the risk of exploited patient information in this phase will decrease by 90%. The reverse is also true. The other minus relationships work in the same way.

We associated an activation threshold for each event that specifies the minimum strength to which the incoming relationship degrees must be aggregated to activate an event. For C_4 (the exploitation of paper files) to occur, the incoming relationships must be aggregated to a minimum of 0.8, or 80%. If the patient spends time in registration (event C_2) and the doctors and nurses in charge shared the patient information (event C_1), then the incoming relationships (e_1,e_4) and (e_2,e_4) need to aggregate to at least 0.8 for the paper files to be exploited (C_4). The thresholds of the other events work in the same way.

Even though the FCM can express these dynamic relationships, it is not possible to create a general FCM that will apply to different health care institutions. There are numerous dynamic aspects that vary from one health care institution to another (as well as from one phase to another in a specific patient route of a specific health care institution)—for example, the technologies the institution uses to store and process patient information. We might be able to identify a few general dynamic aspects that are relevant to any health care institution, but we must treat each institution individually.

Calculating a phase's IT risk value

Most events in a typical health care institution are not easily quantifiable, because they merely constitute vague estimates. A fuzzy-rule-based approach, however, provides a way to model the intuitive fuzzy relationships in more detail (see the FCM in Figure 2). Such an approach constitutes a set of fuzzy rules that converts inputs to output.[6] All the fuzzy rules are fired—that is, are activated—in parallel to some degree.[7] Some of the rules, however, fire to zero degrees, with the result that they will not contribute to the final outcome of the fuzzy-rule-based system.

We can view the FCM edges that we used to represent the fuzzy relationships between events as fuzzy if–then rules. For example, if the strength of security controls in the registration phase increases, then the likelihood of database files containing patient information decrease by a certain degree. Human observation and intuition (which are naturally vague) form the basis for constructing such fuzzy rules.

To illustrate the fuzzy-rule-based approach, we will only consider the registration phase in Figure 1. The principal aim of this approach is to calculate an IT risk value linked to a phase in a specific patient route. This risk value is based on the IT domain a typical patient will be exposed to in a specific phase of his or her route.

After calculated IT risk values for each phase in a patient route, the hospital superintendent can use these IT risk values, for example, to identify those phases in a specific patient route that are critical to information security and then further investigate countermeasures to enhance the hospital's IT security.

As we mentioned earlier, the inputs and the output constitute vague estimates rather than crisp values; such vague estimates define general categories, as opposed to rigid, fixed collections. These categories have more

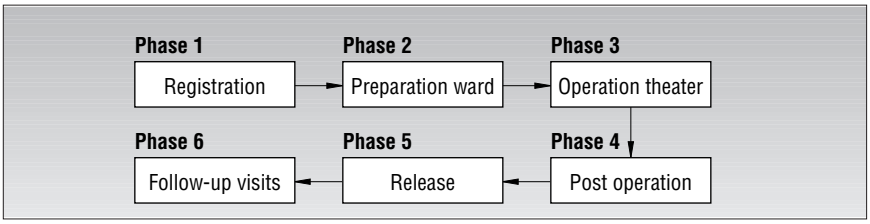

Figure 1. Route followed by a patient admitted to a hospital for an operation.

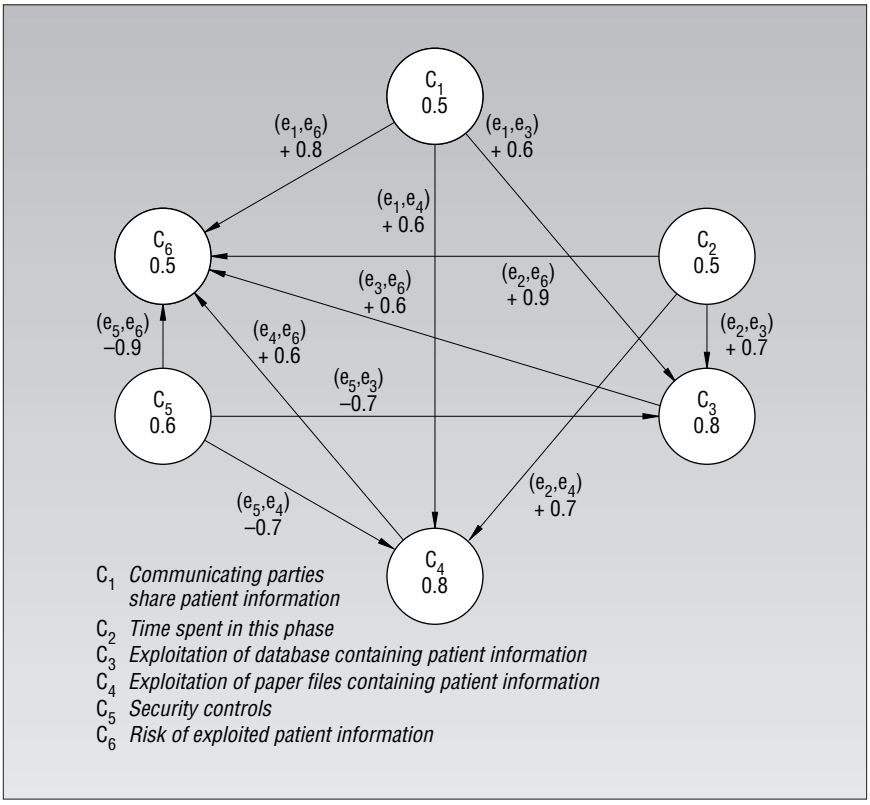

Figure 2. A fuzzy cognitive map, representing the dynamic relationships in a hospital in the registration phase.

flexible membership requirements that allow for partial membership to a category. The degree to which a value is a member of a category can be any value between 0 and 1 (rather than strictly 0 or 1). An estimated value of 11 communicating parties sharing patient information in the registration phase can, for instance, have a membership of 0.8 in the "small number of communicating parties" category. In fuzzy logic, we call such categories *fuzzy sets*.

Each fuzzy set has a corresponding membership function that returns the degree of membership for a given value within a fuzzy set.[7] Figure 3 show how we can represent the inputs and output by means of *membership functions*. Membership functions might take on any form, but the most common shape is a triangle. In Figure 3, we can view the triangular sets as the bisection of a triangular fuzzy set, because they overlap the endpoints of the universe of discourse (or the total allowable values from the smallest to the largest). We use the bell-shaped membership functions in Figure 3 to represent values around a central value.

To convert a series of individual fuzzy regions into a continuous and smooth surface, each fuzzy set a membership function represents must, to some degree, overlap its neighbouring set. This overlap is the natural consequence of fuzziness and ambiguity associated with the segmentation and classification of a continuous space. Experience dictates that the overlap for midpoint-to-edge for neighboring fuzzy regions averages between 25% and 50% of the fuzzy set base. However, drawing membership functions is a matter of common sense and engineering judgment.

Consider Figure 3b: A 420 value for the risk of exploited patient information belongs to the very low fuzzy set to a degree of 0.1, to the low fuzzy set to a degree of 0.35, and to the medium fuzzy set to a degree of 0.82. Therefore, a particular value can belong to more than one fuzzy set at any given time, but the transition from one fuzzy set to the next is gradual.

To determine the influence of the inputs on the risk of exploited patient information (output), we can formulate intuitive linguistic fuzzy rules. We can reason that if the period spent in registration is medium, the security controls in registration are strong, and the number of parties sharing patient information is small or very small, then the risk of exploited patient information is low.

Constructing such fuzzy rules involves intensive knowledge engineering and an understanding of the specific health care institution domain. The aim is to find the best set of rules that reflect the overall behavior of the specific health care domain. This might also involve experimentation with various fuzzy rule sets to determine the rule set that produces the most stable result.

The following fuzzy rules are examples that we can apply to determine the risk value of the registration phase:

IF the number of communicating parties is
very small

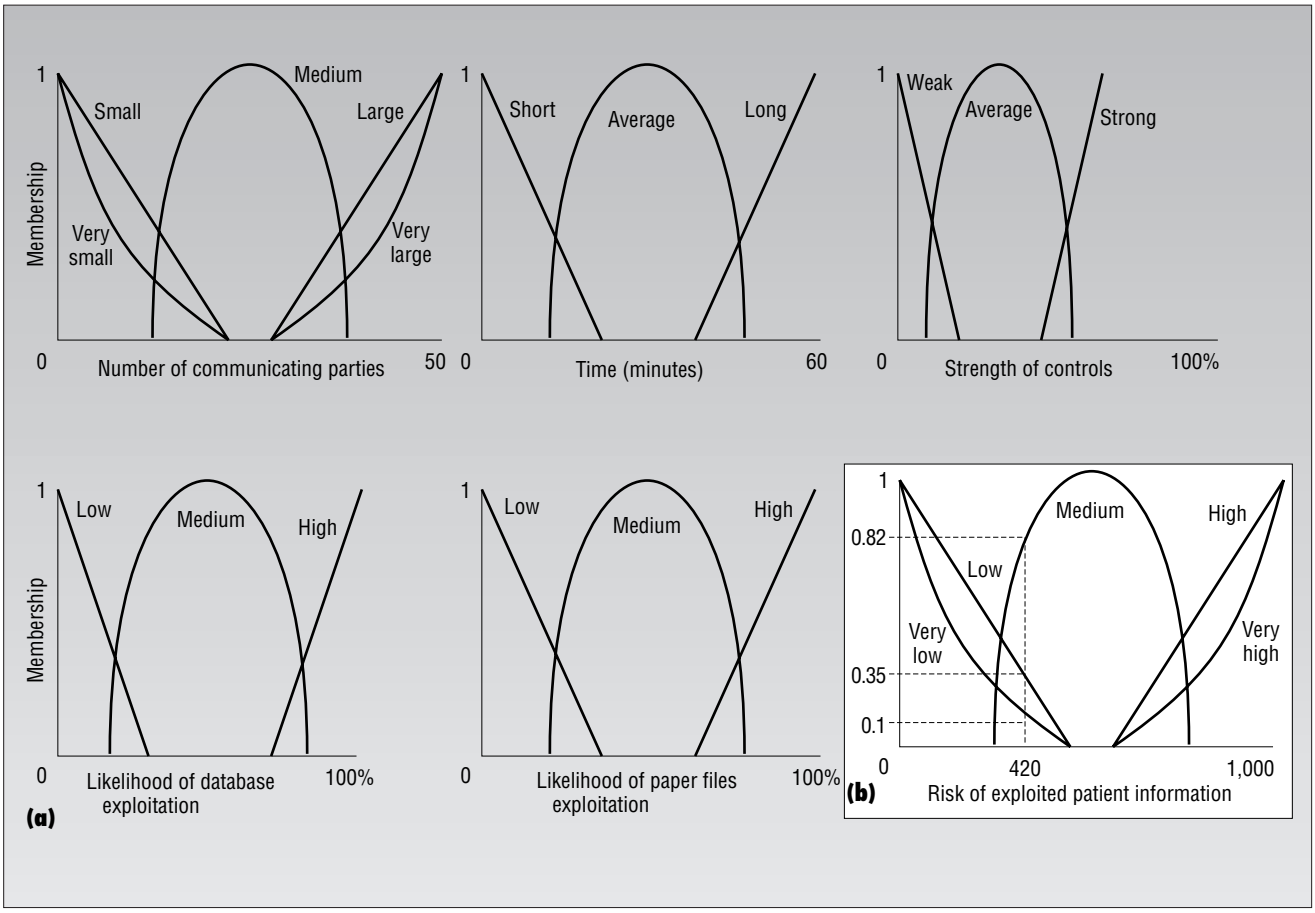

Figure 3. Membership functions for (a) the inputs and (b) the output.

and the time spent in the registration phase is average,
THEN the risk is low.
IF the number of communicating parties is small
and the time spent in the registration phase is average
and the strength of the security controls in place is weak,
THEN the risk is medium.

Consider Figure 4: Supposing that we estimate the number of people sharing the patient information in this phase at 10, then this constitutes a 0.6 membership in the "very small number of communicating parties" fuzzy set and a 0.85 membership in the "small" fuzzy set. Similarly, if we estimate the time the patient spent in the registration phase at 25 minutes and the security controls' strength at 8%, then this constitutes memberships of 0.9 in the "average time spent" fuzzy set and 0.75 in the "weak strength of controls" fuzzy set, respectively.

It might be difficult to distinguish between an estimated strength of 8% and an estimated strength of 12% for a specific security control. In such cases, we should focus on the fuzzy sets as such—for example, we should consider whether the strength of the security control is weak rather than average (so 8% is preferable) or average rather than weak (so 12% is preferable). We do not rigidly specify the exact domain over which we map a specific fuzzy set and the specific shape of the fuzzy set's curve. The initial determination of a fuzzy set's domain and curve shape is normally done intuitively. Fuzzy systems tolerate approximations in the representations of fuzzy sets,[7,8] which means we can possibly use more than one type of fuzzy set to model a specific scenario successfully. After the initial intuitive determination of the curve shapes, we use repeated trial and error system runs to find the optimum configuration for solving a particular problem. According to the scenario in Figure 4, both fuzzy rules we listed earlier will fire to some degree. We must map the input fuzzy sets "very small number of communicating parties" and "average time spent" the first fuzzy rule implied and the input fuzzy sets "small number of communicating parties," "average time spent" and "weak strength of security controls" the second fuzzy rule implied, to the "low" and "medium risk of patient information being exploited" output fuzzy sets, respectively. Figure 5 depicts this process, called "correlation."

Consider the first fuzzy rule—IF the number of communicating parties is very small and the time spent in the registration phase is average, THEN the risk is low. To corre-

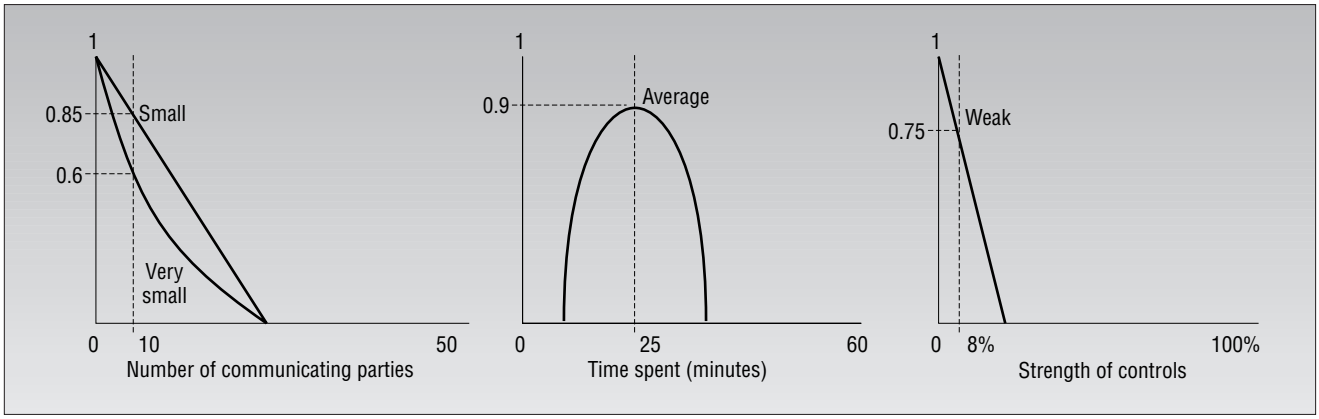

Figure 4. Membership values for the input.

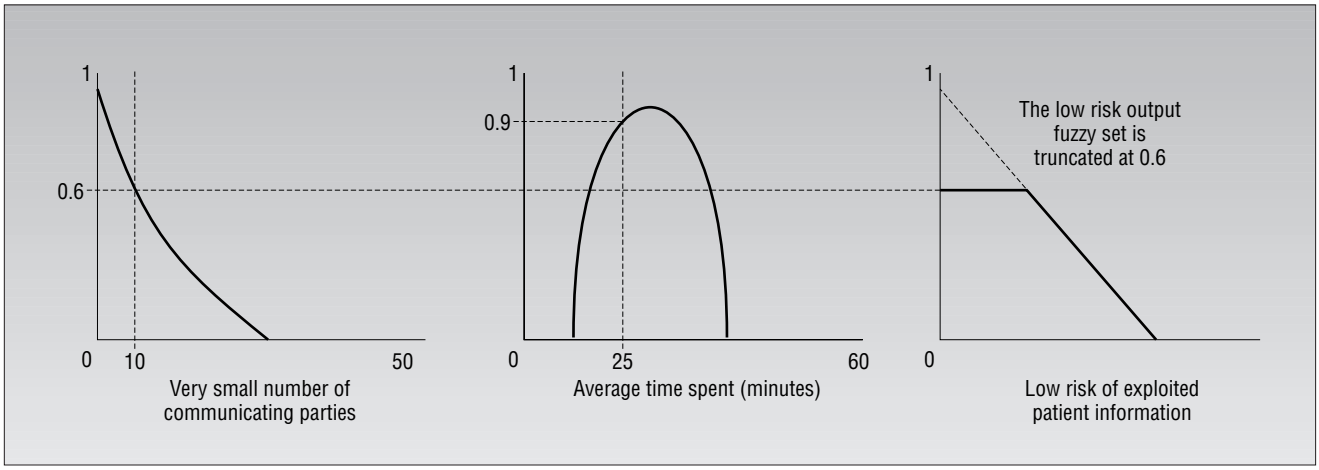

Figure 5. Correlating the input fuzzy sets "very small" and "average" with the "low" output fuzzy set.

late the "very small" and the "average" fuzzy sets with the "low" fuzzy set, we use the correlation minimum method, which truncates the "low risk" fuzzy set at the minimum-truth value of the input fuzzy sets "very small number of communicating parties" and "average time spent." The "very small" fuzzy set's 0.6 membership value is lower than the "average" fuzzy set's 0.9 membership value, so the "low risk" fuzzy set is truncated at 0.6.

The correlation process for the second fuzzy rule—IF the number of communicating parties is small and the time spent in the registration phase is average and the strength of the security controls in place is weak, THEN the risk is medium—works in the same way. The "weak" fuzzy set's 0.75 membership value is lower than the "small" fuzzy set's 0.85 membership value and the "average" fuzzy set's 0.9 membership value (see Figure 4). The "medium risk" fuzzy set is, therefore, truncated at 0.75.

We must aggregate the output fuzzy regions that the two fuzzy rules generate to obtain a combined output fuzzy region. Figure 6 illustrates the aggregation process. The aggregation method we used, namely, the *min/max aggregation method*, takes the maximum of the output fuzzy regions generated at each point along their mutual membership values to produce a final fuzzy region.

The final step in the rule-based approach involves the defuzzification of the output region in a bid to obtain the expected risk value of the registration phase in the patient route under consideration. Figure 6 also illustrates this process.

There are several techniques available for defuzzification. For the purposes of our model, we use the *center of maximum technique* to determine the expected risk value. This technique finds the domain point in the aggregated output region with the maximum truth. The registration phase's risk value is, therefore, 450, which is just below an average risk value on a scale from 0 to 1,000.

Using cognitive fuzzy tools to support decision-making

The fuzzy rule-based approach provided an expected IT risk value linked to the registration phase in the route a typical patient will follow when admitted to hospital for an operation. We can use the FCM introduced earlier to assist management in making decisions based on the outcome of such risk value.

Supposing that the IT risk value for the registration phase is 450, the risk value is lower than average but still not very low. We can use the FCM to explore various what-if scenarios to determine a way to either decrease the risk value or explore whether a scenario could increase the risk value in such a way that the registration phase becomes a high-risk area. What would happen if paper files containing patient information were exploited? We use the FCM to effectively answer such questions.

Consider the relationships between the events in the registration phase the FCM describes in Figure 2. The simple 2D edge matrix in Figure 7 represent these relationships.

The *i*th row lists the connection strength of the edges (e_i,e_k) (which describe relationships) directed out from causal event C_i. The first row in the matrix indicates that the strength of the relationship (e_1,e_3) between C_1 and C_3 is 0.6, the strength of (e_1,e_4) between C_1 and C_4 is 0.6 and that the strength of (e_1,e_6) between C_1 and C_6 is 0.8 (see Figure 2).

Furthermore, C_i causally increases C_k if $(e_i,e_k) > 0$, decreases C_k if $(e_i,e_k) < 0$, and has no effect if $(e_i,e_k) = 0$. Event C_1, for example, causally increases events C_3, C_4, and C_6 to varying degrees, because (e_1,e_3), (e_1,e_4) and (e_1,e_6) are all greater than 0.

Each event in an FCM turns one or more events on (1) or off (0). For example, to model the what-if scenario—namely, what would happen if, for instance, the paper files were exploited in the registration phase—event C_4 needs to be turned on (that is, to be set equal to one). All other events remain zero.

The state vector [0 0 0 1 0 0] can represent this input state. In other words, either a zero or a one in the state vector represents each event in the FCM, depending on whether the vector is turned off or on. Therefore, in our what-if scenario, only the fourth element (C_4) in the state vector has a value of one. FCM input states such as these fire all relationships in the FCM to some degree. This process will show how, in a fuzzy dynamic system, causal events affect each other as time goes by.

To model the effect of the input state I_0 = [0 0 0 1 0 0] (the exploitation of paper files containing patient information) on the FCM for the registration phase, we use the following technique to determine the new state (on or off) for each event C_i each time (t_{n+1}) fires the FCM:

$$c_i(t_{n+1}) = S\left[\sum_{K=1}^{N} e_{ki}(t_n)C_k(t_n)\right]$$

This technique involves a matrix–vector multiplication to transform the weighted input to each event C_i. S(x) is a bounded signal function indicating whether C_i is turned off (0) or on (1).

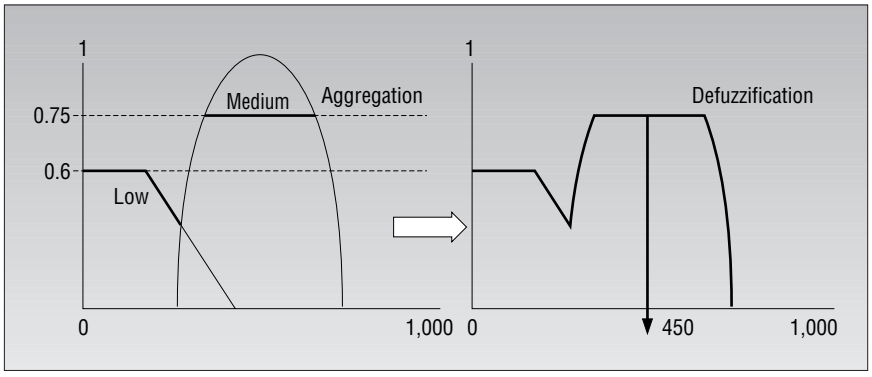

Figure 6. The aggregation and defuzzification process.

E =

	C_1	C_2	C_3	C_4	C_5	C_6
C_1	0	0	0.6	0.6	0	0.8
C_2	0	0	0.7	0.7	0	0.9
C_3	0	0	0	0	0	0.6
C_4	0	0	0	0	0	0.6
C_5	0	0	-0.7	-0.7	0	-0.9
C_6	0	0	0	0	0	0

Figure 7. The edge matrix, E, representing the relationships between the events in the registration phase of the route a patient follows when admitted to hospital for an operation.

We apply the above formula to the FCM with initial input state [0 0 0 1 0 0] as follows

$I_0 = [0\ 0\ 0\ 1\ 0\ 0]$, then

$$I_0E_c = \left[\sum_{k=1}^{6} I_{0k}e_{k1}, \sum_{k=1}^{6} I_{0k}e_{k2}, \sum_{k=1}^{6} I_{0k}e_{k3}, \sum_{k=1}^{6} I_{0k}e_{k4}, \sum_{k=1}^{6} I_{0k}e_{k5}, \sum_{k=1}^{6} I_{0k}e_{k6}\right].$$

I_{0k} refers to the kth element in the state vector $I_0 = [0\ 0\ 0\ 1\ 0\ 0]$

e_{k1} refers to the entry in the kth row in the first column of the edge matrix E

e_{k2} refers to the entry in the kth row in the second column of the edge matrix,

E, and so forth.

$$= [0*0 + 0*0 + 0*0 + 1*0 + 0*0 + 0*0,$$
$$0*0 + 0*0 + 0*0 + 1*0 + 0*0 + 0*0,$$
$$0*0.6 + 0*0.7 + 0*0 + 1*0 + 0*-0.7 + 0*0$$
$$0*0.6 + 0*0.7 + 0*0 + 1*0 + 0*-0.7 + 0*0$$
$$0*0 + 0*0 + 0*0 + 1*0 + 0*0 + 0*0$$
$$0*0.8 + 0*0.9 + 0*0.6 + 1*0.6 + 0*-0.9 + 0*0]$$
$$= [0\ 0\ 0\ 0\ 0\ 0.6]$$
$$\xrightarrow{0.5} I_1 = [0\ 0\ 0\ 1\ 0\ 1]$$

The arrow represents a threshold operation, with 0.5 assumed as the threshold value. In other words, all entries in the state vector I_0E_c with values higher than or equal to 0.5 is turned on. Furthermore, we keep C_4 on, because we want to model the effect of a sustained threat of paper files containing patient information being exploited in the registration phase. Therefore, we can make the following conclusion: When I_0 fires the FCM (or when I_0 occurs), then event C_6 is turned on. The next input state firing the FCM will be $I_1 = [0\ 0\ 0\ 1\ 0\ 1]$.

We can apply the formula to the FCM with input state I_1 in the same way:

$$I_0E_c = \left[\sum_{k=1}^{6} I_{0k}e_{k1}, \sum_{k=1}^{6} I_{0k}e_{k2}, \sum_{k=1}^{6} I_{0k}e_{k3}, \sum_{k=1}^{6} I_{0k}e_{k4}, \sum_{k=1}^{6} I_{0k}e_{k5}, \sum_{k=1}^{6} I_{0k}e_{k6}\right]$$
$$= [0\ 0\ 0\ 0\ 0\ 0.6]$$
$$\xrightarrow{0.5} I_2 = [0\ 0\ 0\ 1\ 0\ 1] = I_1$$

This results in C_6 remaining on. The next input state $I_2 = [0\ 0\ 0\ 1\ 0\ 1]$ is, therefore, equal to the previous input state I_1. The FCM then converges to a fixed point I_2 that turns on C_6, which means that the exploitation of paper files in the registration phase will increase the risk of patient information being exploited (C_6). (A detailed explanation of this technique falls outside the scope of this article: consult Bart Kosko's book, *Fuzzy Engineering*, for more information.[4]) Based on the outcome of the what-if scenario, the institution should control the likelihood of paper files containing patient information being exploited to prevent the registration phase's IT risk value from increasing.

Modeling what-if scenarios with FCMs can greatly assist decision making about security control implementation in a specific phase of a typical patient route. Using this modeling technique, we can identify security services (such as confidentiality, integrity, authentication, authorization, and nonrepudiation) threatened by those specific scenarios and ignore the security services that are not in any danger.

THIS COGNITIVE FUZZY APPROACH is unique because it uses both the FCM and the fuzzy-rule-based techniques to calculate the IT risk value linked to a phase in a specific patient route. The advantage of using these techniques together is that it takes into account intuitive human observation, which forms the basis of any risk assessment, and also accounts for the vagueness regarding patient information and risks when calculating a phase's risk level in a typical patient route. By identifying a phase's IT risk value, this approach helps health care staff manage risks by facilitating the decision-making process.

We've aimed our further research at developing a complete risk-management model specifically tailor-made to suit the health care domain. We will base such a model on the cognitive fuzzy modeling approach we discussed here. Another possibility for further research involves investigating this risk-management model to adapt it to suit other types of organizations. ■

References

1. T.C. Rindfleisch, "Information Technology and Health Care," *Comm. ACM*, Vol. 40, No. 8, Aug. 1997, pp. 93–100.
2. E.D. Cox, *Fuzzy Logic for Business and Industry*, Charles River Media, Rockland, Mass., 1995, p. 601.
3. F.M. McNeill and E. Thro, *Fuzzy Logic: A Practical Approach*, Academic Press, Boston, 1994.
4. B. Kosko, *Fuzzy Engineering*, Prentice Hall, Upper Saddle River, N.J., 1997, p. 549.
5. B. Kosko, "Fuzzy Cognitive Maps," *Int'l J. Man-Machine Studies*, Vol. 24, 1986, pp. 65–75.
6. B. Kosko, *Fuzzy Thinking*, Flamingo, London, 1994, p. 318.
7. E.D. Cox, *The Fuzzy Systems Handbook: A Practitioner's Guide to Building, Using, and Maintaining Fuzzy Systems*, Academic Press, Boston, 1994, p. 515.
8. S.T. Welstead, *Neural Network and Fuzzy Logic Applications in C/C++*, John Wiley & Sons, New York, 1994.

Elmé Smith is a senior lecturer in computer science and information systems at the University of South Africa. Her research interests include IT risk management, health care security, access control, and database security. She received her BSc and MSc in computer science from Potchefstroom University for Christian Higher Education and is currently working toward a PhD in computer security at the Rand Afrikaans University. Contact her at the Computer Science and Information Systems Dept., 8-85 Theo van Wijk building, Muckleneuk, UNISA, Pretoria, South Africa; smithe@alpha.unisa.ac.za.

Jan Eloff is a professor in computer science at Rand Afrikaans University. He gained practical experience by working as a computer management consultant specializing in the field of information security. He is chairman of the Special Interest Group in Information Security and is chairman of the International Working Group 11.2 of IFIP specializing in small systems security. He is an evaluated researcher from the Foundation for Research Development, South Africa. He received a PhD in computer science from Rand Afrikaans University. Contact him at the Computer Science Dept., C-ring 521, Kingsway, Auckland Park, RAU, South Africa; eloff@rkw.rau.ac.za.