Special Issue

JDMS

Journal of Defense Modeling and Simulation: Applications, Methodology, Technology
2017, Vol. 14(1) 17–32

DOI: 10.1177/1548512916680779
journals.sagepub.com/home/dms
SAGE

Fuzzy cognitive maps of public support for insurgency and terrorism

Osonde A Osoba[1,2] and Bart Kosko[3]

Abstract

Feedback fuzzy cognitive maps (FCMs) can model the complex structure of public support for insurgency and terrorism (PSOT). FCMs are fuzzy causal signed digraphs that model degrees of causality in interwoven webs of feedback causality and policy variables. Their nonlinear dynamics permit forward-chaining inference from input causes and policy options to output effects. We show how a concept node causally affects downstream nodes through a weighted product of the intervening causal edge strengths. FCMs allow users to add detailed dynamics and feedback links directly to the causal model. Users can also fuse or combine FCMs from multiple experts by weighting and adding the underlying FCM fuzzy edge matrices. The combined FCM tends to better represent domain knowledge as the expert sample size increases if the expert sample approximates a random sample. Statistical or machine-learning algorithms can use numerical sample data to learn and tune a FCM's causal edges. A differential Hebbian learning law can approximate a PSOT FCM's directed edges of partial causality using time-series training data. The PSOT FCM adapts to the computational factor-tree PSOT model that Davis and OMahony based on prior social science research and case studies. Simulation experiments compare the PSOT models with the adapted FCM models.

Keywords

Fuzzy cognitive maps, causal reasoning, knowledge fusion, differential Hebbian learning

1 Modeling feedback causal webs with fuzzy cognitive maps

This paper presents static and dynamic fuzzy cognitive map (FCM) models of public support for insurgency and terrorism (PSOT). We base these PSOT FCMs on the factor-tree PSOT analysis of Davis et al.[1,2] Public support for insurgency and terrorism has complex sociopolitical causes[3,4] that involve numerous factors. FCMs can efficiently model and process the interwoven causal and policy structure of PSOT and other defense problems. FCM applications number in the thousands and range from control engineering and signal processing to policy analysis and social modeling.[5,6]

FCMs are fuzzy causal signed directed graphs. They are fuzzy because in general both their directed causal edges and their concept nodes are multivalued and so can assume more values than just the extremes of on or off. They locally model degrees of causality through their directed causal edge strengths.[7] Users can express their causal and policy models by drawing signed weighted causal edges between concept nodes. Figure 1 shows a FCM fragment that models an undersea causal web of dolphins in the presence of sharks or other survival threats. The next section shows how to make what-if inferences or predictions with this simple FCM that has binary concept nodes and trivalent causal edges. The inference process uses only vector-matrix multiplication and thresholding. More complex FCMs can activate concept nodes with some of the nonlinear functions in Figure 2 or with many

[1]RAND Corporation, USA
[2]Pardee RAND Graduate School, USA
[3]University of Southern California, USA

Corresponding author:
Osonde A Osoba, The RAND Corporation, 1776 Main St, Santa Monica, CA, 90405, USA.
Email: oosoba@rand.org

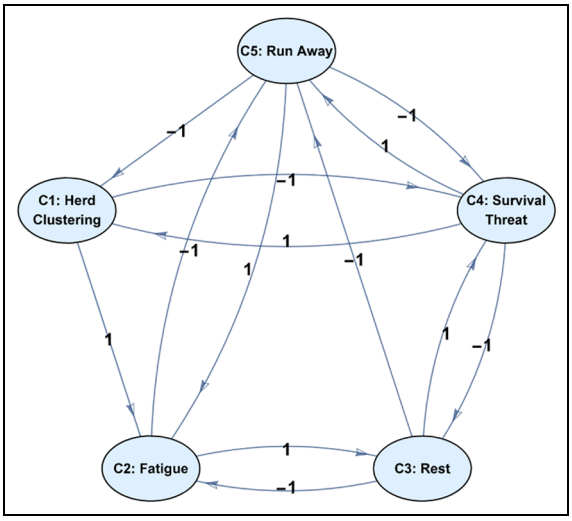

Figure 1. Fragment of a predator–prey fuzzy cognitive map (FCM) that describes dolphin behavior in the presence of sharks or other survival threats.[8] The FCM itself is a fuzzy signed directed graph with feedback. The concept nodes of the digraph represent fuzzy sets that activate to varying degrees of concept-node occurrence. The edges denote fuzzy or partial causal dependence between concept nodes. The edges in this FCM are trivalent: $e_{ij} \in \{-1, 0, 1\}$. Each nonzero edge defines a causal if–then rule. The dolphin pod decreases its resting behavior if a shark or other survival threat is present. But the survival threat increases if the pod rests more. These two causal links define a minimal cycle or feedback loop within the FCM's causal web. Such feedback cycles endow the FCM with transient and equilibrium dynamics. All inputs produce equilibrium limit cycles or fixed-point attractors in the simplest case where all nodes are bivalent threshold functions and when the system updates all nodes at each iteration.

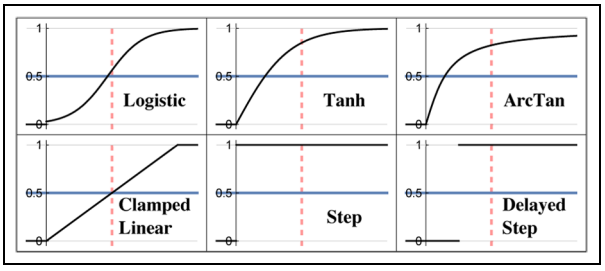

Figure 2. Six types of FCM concept-node occurrence or activation function: sigmoid logistic, sigmoid hyperbolic tangent, arctangent, linear, step function, and the delayed step. Each occurrence function maps into the unit interval $[0, 1]$ and gives the degree to which the concept or policy occurs at a given moment in the causal web. Simulations used the logistic, linear, or step function.

other monotonic or nonmonotonic functions. A causal learning law can approximate the causal edge values given time-series data of the concept nodes. Figure 3 shows the approximation path of one such causal edge from a PSOT FCM.

A FCM's overall cyclic signed digraph structure resembles a feedback neural or semantic network. The graph structure permits inference through forward chaining and allows the user to control the level of causal or conceptual granularity. A FCM concept node can itself be part of another FCM or of some other nonlinear system. Feedforward fuzzy rule-based systems can also model the input-output structure of a concept node just as they can model a single causal edge that connects one concept node to another. Such fuzzy systems are uniform function approximators if they use enough if-then rules.[9] Their rule bases adapt using both unsupervised and supervised learning laws.[10,11]

A FCM's feedback loops model interwoven causal webs and can produce rich and predictive equilibrium dynamics.[12,13] These causal equilibria define "hidden patterns"[12] in the often inscrutable web of edges and nodes. FCMs with binary concept variables produce limit-cycle equilibria or simple fixed-point attractors. Properly fuzzy concept nodes can in principle produce more exotic equilibria such as limit tori or chaotic attractors.

A FCM's underlying matrix structure makes it easy to combine or fuse FCMs from several sources to produce an overall representative FCM. The strong law of large number shows that in many cases this fused FCM converges with probability one to the population FCM of the sampled FCMs.[14] This result holds formally if the FCM values from the combined experts approximate a statistical random sample with finite variance. A random sample is sufficient for this convergence result but not necessary. A combined FCM may still give a representative knowledge base when the expert responses are somewhat correlated or when the experts do not all have the same level of expertise. Users can also construct FCMs from written sources such as policy articles or books or legal testimony. They can also use statistical learning algorithms to grow FCMs from sample data. FCMs can in this way address the growing representational problems of big data and what we have called "big knowledge."[6] Figure 4 shows the minimal fusion case of combining two FCMs with overlapping concept nodes.

Knowledge fusion is a key function in defense and intelligence decision-making processes.[15,16] The PSOT FCMs we develop below allow such knowledge combination or fusion[12,14,17,18] and can use a causal differential Hebbian learning law[12] to grow and tune the FCM causal edges from numerical time-series data.

The FCM PSOT models below rely on the prior PSOT modeling of Davis et al.[1,2] Their carefully constructed PSOT model is a factor tree model. It reflects social science research that describes the factors and societal

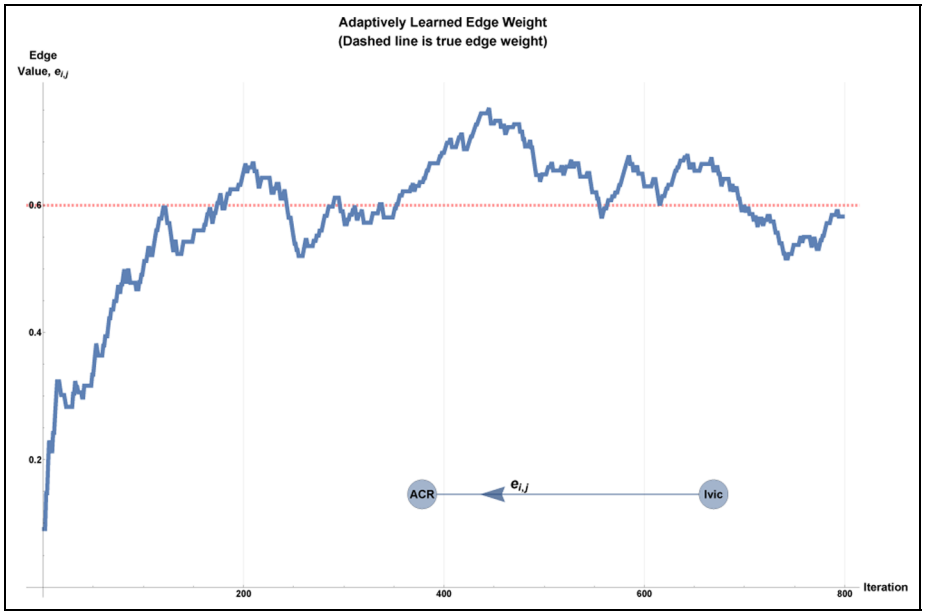

Figure 3. Learning fuzzy cognitive map causal-edge values e_{ij} with time-series data from activated causal concept nodes. The directed causal edge *Ivic* → *ACR* between *Assessment of likely victor* and *Acceptability of costs and risks* is uncertain in the original Public Support for Insurgency and Terrorism model. We can infer the value of this directed causal edge with adaptive inference algorithms such as differential Hebbian learning if we have access to the time-series data history of both concept nodes. The time-series data may come from survey data, field measurements, or expert elicitations. The plot shows that differential Hebbian learning quickly converged to a good approximation of the edge value e_{ij}.

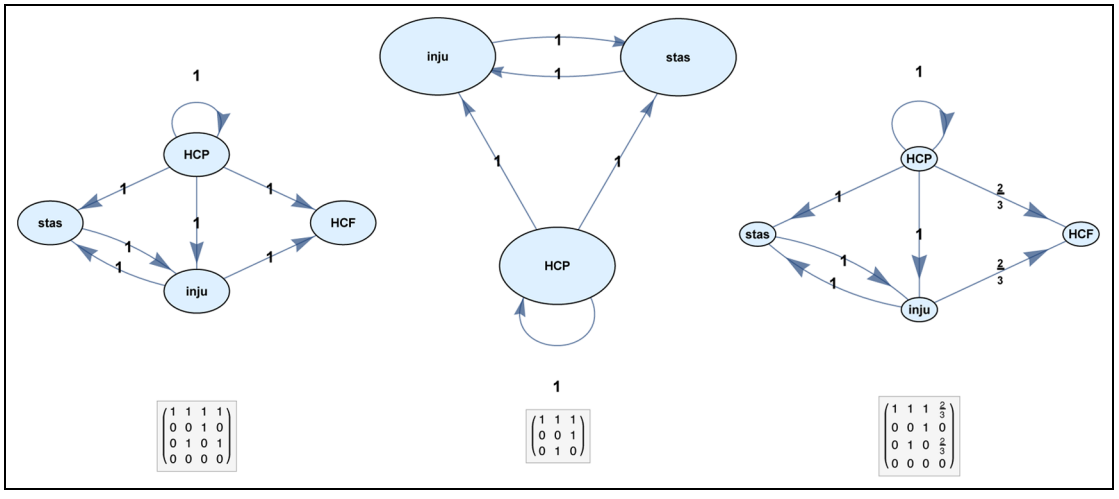

Figure 4. Fuzzy cognitive map (FCM) knowledge combination or fusion by averaging weighted FCM adjacency matrices. The three digraphs show the minimal case of combining two FCMs that have overlapping concept nodes. The third FCM is the weighted combination of the first two. The expert problem domain is the medical problem of strokes and blood clotting involving Virchow's triad: blood stasis (stas), endothelial injury (inju), and hypercoagulation factors (HCP and HCF).[14] Expert 1 has a larger FCM than Expert 2 because Expert 1 uses an extra concept node. We can fuse their knowledge webs by averaging their causal-edge adjacency matrices using Equation (25). This weighted average uses each expert's causal-edge matrix $\mathbf{E} = \frac{2}{3}\mathbf{E}_1 + \frac{1}{3}\mathbf{E}_2$, as shown in the combined (third) FCM. The weighting assumes that the first expert is twice as credible as the second expert. Note that Expert 2 ignored the HCF factor. This results in a new row and column of all zeros in $\mathbf{E}_2$.

relationships that determine a community's propensity to support insurgent or terrorist acts. This PSOT factor tree model has limited scenario simulation capability because the ''public'' under study in the PSOT model is a heterogeneous aggregation. Later work introduced the related propensity for terrorism model[15,16] that focuses on the factors that influence someone's or some group's propensity to support terrorism or carry out terrorist acts.

We use both static and dynamic FCM versions of the PSOT factor tree. The static FCM behaves much as does the original factor tree model. The dynamic FCM extends the static model to allow simulations of long-run what-if scenarios.

The next sections develop the FCM PSOT models. Section 2 explains and extends the basic mathematical and graph structure of FCMs and how to use them for forward-chaining inference in analysis. Section 3 presents the original PSOT model. Then we introduce a FCM version of the PSOT model. This includes adaptations to the PSOT model that allow long-run simulations. We demonstrate the fusion and simulation capabilities of some FCM-PSOT models. Section 4 summarizes the FCM PSOT experiments for causal inference and knowledge fusion.

2 Inference with fuzzy cognitive maps

Fuzzy cognitive maps (FCMs)[6,13] are fuzzy signed directed graphs that describe degrees of causality and webs of causal feedback. Most FCMs have cycles or closed loops that model causal feedback. FCMs can be acyclic and thus define trees. This is rare in practice and implies that such a FCM has no feedback dynamics.

FCMs are fuzzy because their nodes and edges can be multivalent and so need not be binary or bivalent. A property or concept is fuzzy if it admits degrees and is not just black and white.[19–21] Then the property or concept has borders that are gray and not sharp or binary. A subset A of a space X is properly fuzzy if and only if at least one element $x \in X$ belongs to A to a degree other than 0 or 1. Then the set A breaks the so-called ''law'' of contradiction because then $A \cap A^c \neq \emptyset$ holds where A^c is the complement set of A. The set A equivalently breaks the dual ''law'' of excluded middle because then $A \cup A^c \neq X$ holds. Equality holds in these two ''laws'' just in case A is an ordinary bivalent set.

A FCM concept node is fuzzy in general because it can take values in the unit interval [0, 1]. So its values over time define a fuzzy set. This implies that a concept node that describes a survival threat or any other property or policy both occurs and does not occur to some degree at the same time. It cannot both occur 100% and not occur 100% at the same time. The two percentages must sum to 100%. Nor does this preclude applying a probability measure to a concept node or fuzzy set. The probability of a fuzzy event combines the two distinct uncertainty types of randomness and vagueness or fuzz (and formally involves taking the expectation of a measurable fuzzy indicator function[22]). So it makes sense to speak of the probability of a partial survival threat. This differs from the compound uncertainty of a fuzzy probability such as the statement that the survival-threat probability is low or very high. This paper works only with fuzzy concept values.

A directed causal edge e_{ij} is also fuzzy because in general it takes on a continuum of values. The edge can also have a positive or negative sign. So it takes values in the bipolar interval $[-1, 1]$. The use of ''dis-concepts'' can convert all negative causal edges into positive edge values.[7] The edge formally defines a relation or fuzzy subset of a concept product space.

A FCM consists of n concept nodes C_j and n^2 directed fuzzy causal edges e_{ij}. The n concept nodes $C_1, C_2, \ldots, C_n$ are nonlinear and represent variable concepts or factors in a causal system. They are nonlinear in how they convert their inputs to outputs. The concept nodes can define concepts or social patterns that increase or decrease such as political instability or jihadi radicalism. Or they can be policies or control variables that increase or decrease such as weapons spending or foreign investment in a country. The very first FCM published[7] dealt with concepts related to Middle East stability such as Islamic fundamentalism and Soviet imperialism and the strength of the Lebanese government. The author based this first FCM on a 1982 newspaper editorial from political analyst Henry Kissinger titled ''Starting Out in the Direction of Middle East Peace.''[23]

A concept node's occurrence or activation value $C_i(t_k)$ measures the degree to which the concept C_i occurs in the causal web at time t_k. It can also reflect the degree to which it is true that the ith node fires or appears in a given snapshot of the causal web at time t_k. The FCM state vector $\mathbf{C}(t_k)$ gives a snapshot of the FCM system at time t_k.

A FCM model must specify the nonlinear dynamics of the n concept nodes $C_1, C_2, \ldots, C_n$. It must also specify the n^2 directed and signed causal-edge values e_{ij} that connect the ith concept node C_i to C_j. The edges can be time-varying functions in more general FCMs.

We start with the nonlinear structure of the concept nodes. The jth concept node C_j depends at time t_k on a scalar input $x_j(t_k)$ that weights and aggregates all the inflowing causal activation to C_j. Then some nonlinear function Φ_j converts $x_j(t_k)$ into the concept node's new state $C_j(t_{k+1})$ at the next discrete time t_{k+1}. The FCM literature explores discrete and continuous node models with a wide variety of nonlinearities and time lags.[5,6] We present here

the simplest case of a discrete FCM where each node's current state depends on an edge-weighted inner product of the node activity:

$$C_j(t_{k+1}) = \Phi_j\left(\sum_{i=1}^{n} C_i(t_k)e_{ij}(t_k) + I_j(t_k)\right) \quad (1)$$

where $I_j(t_k)$ is some external or exogenous forcing value or input at time t_k. The simplest nonlinear function Φ_j is a hard threshold that produces bivalent or on–off concept-node values:

$$C_j(t_{k+1}) = \begin{cases} 0 & \text{if } \sum_{i=1}^{n} C_i(t_k)\, e_{ij}(t_k) \;+\; I_j(t_k) \le 0 \\ 1 & \text{if } \sum_{i=1}^{n} C_i(t_k)\, e_{ij}(t_k) \;+\; I_j(t_k) > 0 \end{cases} \quad (2)$$

for a zero threshold value.

Fixing the input I_j as some very large positive (or negative) value ensures that C_j stays on (or off) during an inference cycle. We call this ''clamping'' on (or off) the jth concept node C_j. We clamp one or more concept nodes to test a given policy or forcing scenario. Clamping is the only way to drive policy or other nodes that have no causal fan-in from other concept nodes. We show below how to model the sustained presence of a shark in the dolphin FCM of Figure 1 by clamping on the fourth concept node.

Continuous-valued concept nodes often use a monotonic increasing Φ_j nonlinearity such as the logistic sigmoid function. But Φ_j can also be nonmonotonic. This happens if it is a Gaussian or Cauchy probability density function. It can also be multimodal by forming a mixture of such unimodal probability curves. Then the Expectation-Maximization algorithm can tune the mixture parameters based on numerical training data.[24] Almost all concept nodes are monotonically nondecreasing in the FCM literature. The causal-influence theorem below holds for such activation functions Φ_j.

The logistic causal activation gives a soft threshold that approximates the hard threshold in Equation (2) if the shape parameter $c > 0$ is large enough:

$$C_j(t_{k+1}) = \frac{1}{1 + \exp\left(-c\sum_{i=1}^{n} C_i(t_k)e_{ij}(t_k) - cI_j(t_k)\right)}. \quad (3)$$

The first graph in Figure 2 shows a logistic function and its sigmoidal or soft-threshold shape. The other five graphs show other common causal activation functions. Logistic units are popular in causal and neural-network learning algorithms because they smoothly approximate the on-off behavior of threshold units and still have a simple partial derivative of the form

$$\frac{\partial C_j(x)}{\partial x} = cC_j(1 - C_j) > 0 \quad (4)$$

if

$$C_j(x) = \frac{1}{1 + \exp(-cx)} \quad (5)$$

for scaling constant $c > 0$. The positive derivative in Equation (5) greatly simplifies many learning algorithms. We will also use it below to show the transitive product effect of the edges e_{ij} in a causal inference.

We turn next to the causal edge values e_{ij}. These values are constants during most FCM inferences. The last section below shows how a version of the differential Hebbian learning law can learn and tune these causal edge values from time-series data.

The causal edge value $e_{ij}(t_k)$ in Equation (1) measures the degree that concept node C_i causes concept node C_j at time t_k:

$$e_{ij} = \text{Degree}\left(C_i \to C_j\right). \quad (6)$$

These n^2 causal edge values define the FCM's $n \times n$ fuzzy adjacency matrix or causal edge matrix $\mathbf{E}$. The ith row lists the causal-edge values $e_{i1}, e_{i2}, \ldots, e_{in}$ that flow out from C_i to the other concept nodes (including to itself). The jth column lists the causal-edge values $e_{1j}, e_{2j}, \ldots, e_{nj}$ that flow into C_j from the other concept nodes. So the ith row defines the causal *fan-out* vector of concept node C_i. The jth column defines the causal *fan-in* vector of C_j. The matrix diagonal lists any causal self-excitation of the n concept nodes.

We can also interpret e_{ij} in terms of fuzzy subsethood.[25,26] Then e_{ij} states the degree to which the fuzzy concept set C_j is a fuzzy or partial subset of fuzzy concept set C_j.[7] This abstract framework implies that the edge value e_{ij} is the degree to which the fuzzy concept set C_i belongs to the fuzzy power set of fuzzy set C_j.

A probabilistic view might interpret the edge value e_{ij} as the conditional probability $P(C_j|C_i)$ that C_j occurs given that C_i occurs. An immediate problem is that e_{ij} takes on negative values in the bipolar interval $[-1, 1]$ to indicate causal decrease. There is a simple but somewhat costly way to address this. The original FCM paper[7] showed how to introduce n companion *dis-concepts* to keep all causal-edge values nonnegative and thus how to convert causal decrease into causal increase: ''Extreme terrorism decreases government stability'' holds just in case ''Extreme terrorism increases government instability'' holds. So dis-concepts negate the noun and not the adjective that modifies it. Using dis-concepts doubles the number of concept nodes and expands the edge matrix E to a $2n \times 2n$ matrix. The technique does preserve more causal structure when combining multiple FCMs because then two combined edges of opposite polarity but the same magnitude do not cancel each other out.

There are two structural problems with viewing the directed (positive) edge e_{ij} as the conditional probability

$P(C_j|C_i)$. The first problem is that conditional probability is not transitive but causal implication is transitive. The transitive equality $P(C|A) = P(B|A)P(C|B)$ does not hold in general. A simple counter-example takes any two disjoint or mutually exclusive events A and B with positive joint probabilities $P(A \cap C) > 0$ and $P(B \cap C) > 0$ if all three set events have positive probability. Then $P(C|A) > 0$ but $P(B|A)P(C|B) = 0$ because $P(B|A) = P(A \cap B)/P(A) = 0$ since $A \cap B = \emptyset$. An even starker counter-example results if $A \subset C$ because then $P(C|A) = 1$ while $P(B|A)P(C|B) = 0$.

The second and deeper problem with a probability interpretation of e_{ij} is that it collides with the Lewis impossibility theorem.[27,28] This triviality result and its progeny show that we cannot in general equate the probability of the logical if-then conditional $A \rightarrow B$ with the conditional probability $P(B|A)$. The equality $P(A \rightarrow B) = P(B|A)$ holds only in the trivial case when A and B are independent and thus when there is no conditional relationship at all. So a probabilistic transitive equality of the form $P(A \rightarrow C) = P(A \rightarrow B)P(B \rightarrow C)$ lacks a formal foundation in general. One approach is to replace conditional probability with a more general probable equivalence relation. This gives upper and lower conditional probabilities based on the general inequality that $P(A)P(B) \geq P(A \cap B)P(A \cup B)$[26] because then $P(B|A) \leq Q(B|A)$ if $Q(B|A) = P(B)/P(A \cup B)$. But the resulting conditioning interval does not directly address the basic prohibition that lies behind Lewis's triviality theorem. So a meta-level heuristic may be the best we can make of probabilistic interpretations of the directed edge e_{ij}. Such interpretations may be intuitive but they remain only heuristics.

We now show that FCM nodes influence one another through a weighted product of intervening causal-edge strengths e_{ij}. This result describes a type of causal chaining along a directed path or summed over all such direct paths that connect two concept nodes. It depends on the transitive causal product $e_{j_1 j_2} e_{j_2 j_3} \cdots e_{j_k j_{k+1}}$.

Consider first the directed causal path from concept node C_i to node C_k by way of the intervening node C_j:

$$C_i \underset{e_{ij}}{\longrightarrow} C_j \underset{e_{jk}}{\longrightarrow} C_k \tag{7}$$

Then how does a change in the input node C_i causally affect the downstream node C_k? The chain rule of differential calculus gives a transitive-based product answer for the logistic concept-node activation in Equation (3):

$$\frac{\partial C_k}{\partial C_i} = \frac{\partial C_k}{\partial C_j}\frac{\partial C_j}{\partial C_i} \tag{8}$$

$$= \frac{\partial C_k}{\partial x_k}\frac{\partial x_k}{\partial C_j}\frac{\partial C_j}{\partial x_j}\frac{\partial x_j}{\partial C_i} \tag{9}$$

$$= C_k(x_k)(1 - C_k(x_k))e_{jk}C_j(x_j)(1 - C_j(x_j))e_{ij} \tag{10}$$

$$= e_{ij}e_{jk}\psi_{j,k} \tag{11}$$

using Equations (4) and (5) if we define $\psi_{j,k} = C_j C_k(1 - C_j)(1 - C_k)$. The weighting function $\psi_{j,k} \geq 0$ is maximal when $C_j = 1 - C_j = \frac{1}{2} = C_k = 1 - C_k$ holds for the fuzzy concept nodes C_j and C_k.

So the induced causal effect of a change in C_i depends directly on the causal-edge product $e_{ij}e_{jk}$. This causal influence decays in intensity the lesser C_j or C_k fires or occurs. The edge product $e_{ij}e_{jk}$ is negative if exactly one of the edge values is negative. It is positive otherwise.

The causal-influence result of Equation (11) extends directly to longer causal chains. Suppose there is a directed causal path of length k from the initial concept node C_{j_1} to the final node $C_{j_{k+1}}$:

$$C_{j_1} \underset{e_{j_1 j_2}}{\longrightarrow} C_{j_2} \underset{e_{j_2 j_3}}{\longrightarrow} \cdots \underset{e_{j_k j_{k+1}}}{\longrightarrow} C_{j_{k+1}} \tag{12}$$

Then the chain rule and Equations (4) and (5) again give the influence of C_{j_1} on C_{j_k} as a weighted product of the intervening causal-edge strengths:

$$\frac{\partial C_{j_{k+1}}}{\partial C_{j_1}} = \prod_{l=1}^{k} e_{j_l j_{l+1}} \psi_{j_2, j_3, \ldots, j_{k+1}} \tag{13}$$

where now the nonnegative weighting function $\psi_{j_2, j_3, \ldots, j_{k+1}}$ is the double product $\psi_{j_2, j_3, \ldots, j_{k+1}} = \prod_{l=2}^{k+1} C_{j_l} \prod_{l=2}^{k+1} (1 - C_{j_l})$. The edge product $\prod_{l=1}^{k} e_{j_l j_{l+1}}$ is positive if the number of negative edges is even. It is negative if the number of negative edges is odd. The magnitude of the change $\partial C_{j_{k+1}}/\partial C_{j_1}$ can only decrease as the causal chain lengthens. The fuzziness or partial firing of the concept nodes only exacerbates this monotone causal decay.

The causal influence in Equation (13) still holds if we replace the logistic activation function of Equation (1) of concept node C_j with an arbitrary monotonically nondecreasing functions Φ_j. Then $\partial C_{j_l}/\partial x_{j_l} \geq 0$ and so $\psi_{j_2, j_3, \ldots, j_{k+1}} \geq 0$ because the weighting function is just the product of these activation partial derivatives. This general result on FCM causal influence is important enough to state as a theorem.

Theorem 1. Partial causal influence in fuzzy cognitive maps. *Suppose a fuzzy cognitive map has n concept nodes C_j and n^2 directed causal edges e_{ij}. Suppose further that the concept nodes have monotonically nondecreasing activations: $\partial C_j/\partial x_j \geq 0$ where the argument x_j of $C_j(x_j)$ has the same inner-product form as in Equation (1). Then the causal influence of the concept node C_{j_1} on the downstream node C_{j_k} of the length-k directed causal chain*

$$C_{j_1} \underset{e_{j_1 j_2}}{\longrightarrow} C_{j_2} \underset{e_{j_2 j_3}}{\longrightarrow} \cdots \underset{e_{j_k j_{k+1}}}{\longrightarrow} C_{j_{k+1}} \tag{14}$$

is a nonnegatively weighted product of the intervening causal-edge strengths $e_{j_1 j_2}, \ldots, e_{j_k j_{k+1}}$:

$$\frac{\partial C_{j_{k+1}}}{\partial C_{j_1}} = \prod_{l=1}^{k} e_{j_l j_{l+1}} \psi_{j_2, j_3, \ldots, j_{k+1}} \tag{15}$$

where the weighting function $\psi_{j_2, j_3, \ldots, j_{k+1}}$ *has the form*:

$$\psi_{j_2, j_3, \ldots, j_{k+1}} = \prod_{l=2}^{k+1} \frac{\partial C_{j_l}}{\partial x_{j_l}}. \tag{16}$$

Proof. The result follows from iterated applications of the chain rule:

$$\frac{\partial C_{j_{k+1}}}{\partial C_{j_1}} = \frac{\partial C_{j_{k+1}}}{\partial C_{j_k}} \frac{\partial C_{j_k}}{\partial C_{j_{k-1}}} \cdots \frac{\partial C_{j_2}}{\partial C_{j_1}} \tag{17}$$

$$= \frac{\partial C_{j_{k+1}}}{\partial x_{j_{k+1}}} \frac{\partial x_{j_{k+1}}}{\partial C_{j_k}} \frac{\partial C_{j_k}}{\partial x_{j_k}} \frac{\partial x_{j_k}}{\partial C_{j_{k-1}}} \cdots \frac{\partial C_{j_2}}{\partial x_{j_2}} \frac{\partial x_{j_2}}{\partial C_{j_1}} \tag{18}$$

$$= \frac{\partial C_{j_{k+1}}}{\partial x_{j_{k+1}}} e_{j_k j_{k+1}} \frac{\partial C_{j_k}}{\partial x_{j_k}} e_{j_{k-1} j_k} \cdots \frac{\partial C_{j_2}}{\partial x_{j_2}} e_{12} \tag{19}$$

$$= \prod_{l=1}^{k} e_{j_l j_{l+1}} \prod_{l=2}^{k+1} \frac{\partial C_{j_l}}{\partial x_{j_l}} \tag{20}$$

$$= \prod_{l=1}^{k} e_{j_l j_{l+1}} \psi_{j_2, j_3, \ldots, j_{k+1}} \tag{21}$$

The theorem states only a partial causal result for just one causal path from concept node C_{j_1} to $C_{j_{k+1}}$. A FCM may contain many other directed causal paths from C_{j_1} to $C_{j_{k+1}}$. So the *total* causal change $\mathrm{d}C_{j_{k+1}}/\mathrm{d}C_{j_1}$ invokes the more general chain rule that sums over all the partial derivatives in Equation (17) in all the paths involved. A discrete version of the theorem also holds. It requires keeping track of the discrete time steps as the causal activation flows from one node in the path to the next node.

We next develop a simple example of FCM inference. This example shows how a FCM answers a policy-based what-if question by converging to a limit-cycle equilibrium. The limit cycle itself is the policy answer.

Consider again the FCM fragment in Figure 1 that describes some of the predator-prey behavior of a dolphin pod in the presence of sharks or other survival threats.[8] The concept nodes are binary with threshold activations that obey Equation (2). Bivalent nodes simplify the dynamical analysis because updating all n nodes at the same time must lead to either a fixed-point attractor or a limit-cycle of bit vectors.

The edges in the dolphin FCM fragment in Figure 1 are trivalent: $e_{ij} \in \{-1, 0, 1\}$. So an edge describes maximal causal increase ($e_{ij} = 1$) or maximal causal decrease ($e_{ij} = -1$) or there is no causal relationship at all ($e_{ij} = 0$). The causal edge adjacency matrix $\mathbf{E}$ for the FCM in Figure 1 is a five-by-five trivalent matrix:

$$\mathbf{E} = \begin{array}{c} \\ C_1 \\ C_2 \\ C_3 \\ C_4 \\ C_5 \end{array} \begin{array}{c} \begin{array}{ccccc} C_1 & C_2 & C_3 & C_4 & C_5 \end{array} \\ \begin{pmatrix} 0 & 1 & 0 & -1 & 0 \\ 0 & 0 & 1 & 0 & -1 \\ 0 & -1 & 0 & 1 & -1 \\ 1 & 0 & -1 & 0 & 1 \\ -1 & 1 & 0 & -1 & 0 \end{pmatrix} \end{array} \tag{22}$$

A key argument for using trivalent edge weights e_{ij} in $\{-1, 0, 1\}$ here and elsewhere is that experts may find it hard to accurately state a graded measure of causal intensity $e_{ij} \in [-1, 1]$ for a causal dependence. It is usually much easier to elicit just sign values from experts than real-valued magnitudes. Taber et al.[14] refer to this difficulty as the expert's *articulation burden*. Real-valued magnitudes also tend to be less reliable. Experts are far more likely to agree on edge signs than on both signs and magnitudes. Even the same expert may state different edge-value magnitudes at different times. This articulation burden motivates averaging the trivalent-edge-valued FCMs of experts to approximate the unknown population FCM.

The stochastic convergence result in the appendix of[14] shows that averaging FCMs with trivalent edges approximates the underlying population FCM that has real edge values. FCM sample averages converge with probability one to the population average in accord with the strong law of large numbers. The underlying limit-cycle structure of the averaged FCM also appears to approximate the limit-cycle structure of the original or population FCM if the concept nodes are binary. The limit-cycle results in[14] are only preliminary simulations. So far no theoretical guarantee of limit-cycle convergence has appeared in the FCM literature.

FCM dynamics depend on the FCM's nonlinear feedback structure. The long-run evolution of the FCM state vector $\mathbf{C}$:

$$\lim_{t \to \infty} \mathbf{C}(t) \tag{23}$$

depends on the initial state $\mathbf{C}(0)$ as well as on the nonlinear structure of the concept nodes and the structure of the FCM causal edge matrix $\mathbf{E}$. Simple two-state or binary-node FCMs converge either to a fixed-point attractor

$$\mathbf{C}^*(t+1) = \Phi(\mathbf{C}^*(t)\ \mathbf{E}) \tag{24}$$

or to a limit cycle of repeating bit vectors. This convergence assumes synchronous updating of all the concept nodes at each time step. This stability or convergence guarantee for binary-node FCMs follows from the general result that every square connection matrix is temporally stable.[13,29]

We now show how a limit-cycle hidden pattern occurs in the dolphin FCM in Figure 1. Suppose that a shark appears at time $t = 0$. Then the fourth or survival-threat concept node occurs or turns on. We can represent this initial state $\mathbf{C}(0)$ of the FCM with the unit bit vector

$$\mathbf{C}(0) = (0, 0, 0, 1, 0).$$

Each of the 5 concept nodes acts as a threshold function with zero threshold as in Equation (2). So $C_k(t) = 1$ if and only if its total inner-product input x is positive: $x > 0$. It otherwise equals zero and thus turns off or stays off if it is not active. Then a forward inference gives the following sequence of FCM state vectors:

$$\begin{aligned}
&\mathbf{C}(0)\mathbf{E} = (1, 0, -1, 0, 1) && \rightarrow (1, 0, 0, 0, 1) = \mathbf{C}(1)\\
&\mathbf{C}(1)\mathbf{E} = (-1, 2, 0, -2, 0) && \rightarrow (0, 1, 0, 0, 0) = \mathbf{C}(2)\\
&\mathbf{C}(2)\mathbf{E} = (0, 0, 1, 0, -1) && \rightarrow (0, 0, 1, 0, 0) = \mathbf{C}(3)\\
&\mathbf{C}(3)\mathbf{E} = (0, -1, 0, 1, -1) && \rightarrow (0, 0, 0, 1, 0) = \mathbf{C}(0)
\end{aligned}$$

This inference sequence defines an equilibrium 4-step limit cycle because the fourth state vector $\mathbf{C}(4) = (0, 0, 0, 1, 0)$ is just the first state vector $\mathbf{C}(0)$. So the FCM equilibrium or hidden pattern is the indefinitely repeating cycle $\mathbf{C}(0) \rightarrow \mathbf{C}(1) \rightarrow \mathbf{C}(2) \rightarrow \mathbf{C}(3) \rightarrow \mathbf{C}(0) \rightarrow \cdots$. This cycle defines the equivalent cycle of bit vectors $(0, 0, 0, 1, 0) \rightarrow (1, 0, 0, 0, 1) \rightarrow (0, 1, 0, 0, 0) \rightarrow (0, 0, 1, 0, 0) \rightarrow (0, 0, 0, 1, 0) \rightarrow \cdots$. The repeating cycle predicts a predator-prey oscillation: The shark threat appears. Then the threatened dolphin pod clusters and runs away. Then the dolphins get tired. Then they rest. But the resting dolphins then attract a shark and so on. This limit cycle can model an incidental appearance of a shark.

Suppose instead that a shark appears and actively pursues the dolphins. We can model this what-if policy scenario by clamping the fourth node on during each update. This again amounts to adding a large positive input value for I_4 in Equation (1). Clamping leads to two transient bit-vector states and then a stable 3-step equilibrium limit cycle:

$$\begin{aligned}
&\mathbf{C}(0)\mathbf{E} = (1, 0, -1, 0, 1) && \rightarrow (1, 0, 0, 1, 1) = \mathbf{C}(1)\\
&\text{since we keep } C_4 = 1 \text{ throughout}\\
&\mathbf{C}(1)\mathbf{E} = (0, 2, -1, -2, 1) && \rightarrow (0, 1, 0, 1, 1) = \mathbf{C}(2)\\
&\mathbf{C}(2)\mathbf{E} = (0, 1, 0, -1, 0) && \rightarrow (0, 1, 0, 1, 0) = \mathbf{C}(3)\\
&\mathbf{C}(3)\mathbf{E} = (1, 0, 0, 0, 0) && \rightarrow (1, 0, 0, 1, 0) = \mathbf{C}(4)\\
&\mathbf{C}(4)\mathbf{E} = (1, 1, -1, -1, 1) && \rightarrow (1, 1, 0, 1, 1) = \mathbf{C}(5)\\
&\mathbf{C}(5)\mathbf{E} = (0, 2, 0, -2, 0) && \rightarrow (0, 1, 0, 1, 0) = \mathbf{C}(3)
\end{aligned}$$

The equilibrium three-step limit cycle is $\mathbf{C}(3) \rightarrow \mathbf{C}(4) \rightarrow \mathbf{C}(5) \rightarrow \mathbf{C}(3) \rightarrow \cdots$ or $(0, 1, 0, 1, 0) \rightarrow (1, 0, 0, 1, 0) \rightarrow (1, 1, 0, 1, 1) \rightarrow (0, 1, 0, 1, 0) \rightarrow \cdots$. The limit cycle defines and thus predicts a different form of predator–prey behavior. The shark tires the dolphin pod. The dolphins cluster in a safety maneuver. They then try to rest and still run away as they fatigue. The shark does not relent and the dolphins fatigue and so on.

We show next how FCM models naturally combine or fuse knowledge networks from multiple experts. A group of m experts can each produce an FCM causal edge matrix $\mathbf{E}_k$ that describes their understanding of the prey system in Figure 1. A simple and powerful way to fuse these expert opinions is to take the weighted average of the panel's knowledge base or FCMs by taking the convex combination of their edge matrices[12,13,18]:

$$\bar{\mathbf{E}}_m = \sum_{k=1}^{m} w_k \mathbf{E}_k \tag{25}$$

where the weights w_k are convex weights and hence nonnegative and sum to one as in Figure 4.

The weights w_k can reflect relative expert credibility in the problem domain. So the weights can correspond to test scores or to subjective valuations or to some other measure of the experts' predictive accuracy in prior experiments. Predd et al.[30] developed a method for aggregating expert contributions in cases where experts can abstain or be incoherent. We simply take the weights as given and use equal weights as a default.

The edge matrices $\mathbf{E}_k$ in Equation (25) must be conformable for addition. So they must have the same number of rows and columns and they must be in the same matrix positions. We take the union of all concept nodes from all m knowledge sources. This gives a total of n distinct concept nodes. We zero-pad or add rows and columns of zeros for missing nodes in a given knowledge source's causal edge matrix. This produces a conformable n-by-n adjacency matrix $\mathbf{E}_k$ after appropriately permuting rows and columns to bring them in mutual coincidence with all other zero-padded augmented matrices.

The strong law of large numbers gives some guarantees about the convergence of this fusion knowledge graph to a representative population FCM if the knowledge sources are approximately statistically independent and identically distributed and if they have finite variance.[12,14] Then the weighted average in Equation (25) can only reduce the inherent variance in the expert sample FCMs. So the knowledge fusion process improves with sample size m.

FCMs are digraph models and so resemble Bayesian belief networks (BBNs) and decision trees and factor trees. BBNs are directed *acyclic* probabilistic graphs that represent causal dependence among random variables. They form the basis of Pearl's alternate model of causal inference.[31]

FCMs differ from BBNs in many respects. They differ conceptually because an FCM's nodes need not represent random variables and in practice seldom do. Partial causality is causality that occurs only to some degree. It is not a probability or bet that the cause or effect occurs all or none. More complex FCMs can superimpose such randomness on top of the fuzzy degrees of occurrence. Even then the two types of uncertainty are distinct even though they merge and produce a single real number.

FCMs differ dynamically from BBNs because most FCMs are rich nonlinear dynamical systems while BBNs are feedforward trees and have no dynamical structure. FCM dynamics and inference depend on their nonlinear concept-node functions and on their edge-based feedback cycles. More complex FCMs replace the constant edges of ordinary FCMs with nonlinear functions that vary with time. This leads to still more complex transient and equilibrium dynamics. The lack of cycles in BBNs precludes any nontrivial dynamics. But a BBN's acyclic structure may permit finer control when propagating probabilistic beliefs. Probabilistic belief propagation is also NP-hard.[32] This complexity can impose a heavy computational burden for large BBNs.

The cycles in a FCM directly model feedback causality among the concept nodes. This cyclic structure gives rise in turn to complex dynamics that range from simpled fixed-point attractors and limit cycles to chaotic or aperiodic attractors in more advanced FCMs. The dynamic attractor regions partition the FCM's state space into a finite number of such regions. Every input state converges to exactly one of these regions. The mapping from inputs to regions serves as a macro form of stored input-output associations or what-if questions and answers.

FCMs also differ from BBNs in how they combine expert knowledge sources. Trees do not naturally combine to produce a tree because cycles tend to appear among the nodes. So m BBN probability trees do not naturally combine to form a representative BBN. So such knowledge fusion does not improve with the expert or knowledge-engineer sample size m. But FCMs always combine to yield a new FCM from the matrix averaging process Equation (25). So FCMs are closed under knowledge combination while BBNs are not. The same holds for AI search trees or any other knowledge representation structure based on acylic graphs.

3 The PSOT model: public support for insurgency and terrorism

The Public Support for Insurgency and Terrorism (PSOT) model[1,2] is a factor-tree model that Davis[33] developed to describe the factors and causal pathways that influence a public's support for insurgent or terrorist organizations and actions. The PSOT model synthesizes prior social-science research on terrorism and social movements theory.[2,34,35] This work has validated the PSOT model on case studies of terrorist groups. These groups include al-Qa'ida and the Taliban in Afghanistan, the Kurdistan Workers' Party in Turkey, and the Maoists in Nepal. More recent work[1] has distilled the extensive prior social science research on the topic into a computational PSOT model. Davis's later work[15,16] used the PSOT model to motivate related models of an individual propensity for terrorism.

The PSOT model is a causal factor tree model because it depicts the degree to which child nodes influence or cause parent nodes. Figure 5 and Table I give more details on the components and structure of the PSOT factor tree. The PSOT nodes represent factors that directly or indirectly relate to the Public Support for Insurgency and Terrorism concept PSOT.

Davis's factor tree models are multi-resolution models.[36] Major elements have a hierarchical structure that allows users to specify factors at different levels of detail. Each node is an exogenously driven factor or it fires or activates based on a function of its inputs.

There are also cross-cutting factors besides sub-node factors. Cross-cutting factors affect multiple factors simultaneously. The "*~and*" nodes depend on all fan-in factors being present to a first approximation. The "*~or*" nodes depend on any of the fan-in factors being active or on a combination of the fan-in factors being active. There are several top-level factors that directly relate to the general PSOT of Davis et al.[2]: effectiveness of the organization *EFF*, motivation for supporting the group or cause *MOTV*, the perceived legitimacy of violence *PLEG*, and the acceptability of costs and risks *ACR*. Each of these factors has attendant contributory sub-factors.

PSOT edges denote positive influences by default. We denote negative edges with '$-$' as with a FCM causal-decrease edge. Factor activation along a negative edge reduces the activation of the parent factor. We denote *ambiguous* edges with "$+/-$". The ambiguity refers to uncertainty over the edge's direction of influence.

We based our FCM models on the important case of the al-Qa'ida transnational terrorist organization.

Davis et al.[2] have discussed how the PSOT model explains public support for al-Qa'ida's mission as follows (paraphrased from Davis et al.[2]). The organizational effectiveness of al-Qa'ida depends in part on the charisma, strategic thinking, and organizational skills of its leadership (*lead*). al-Qa'ida has packaged and framed its ideology to appeal to many Muslims worldwide. Motivation for public support of al-Qa'ida's beliefs comes from shared religious beliefs that stress common identity (*id*) and the sense of duty (*duty*) that such identity fosters. al-Qa'ida also relies on a popular narrative of shared grievances (*shgr*) in the

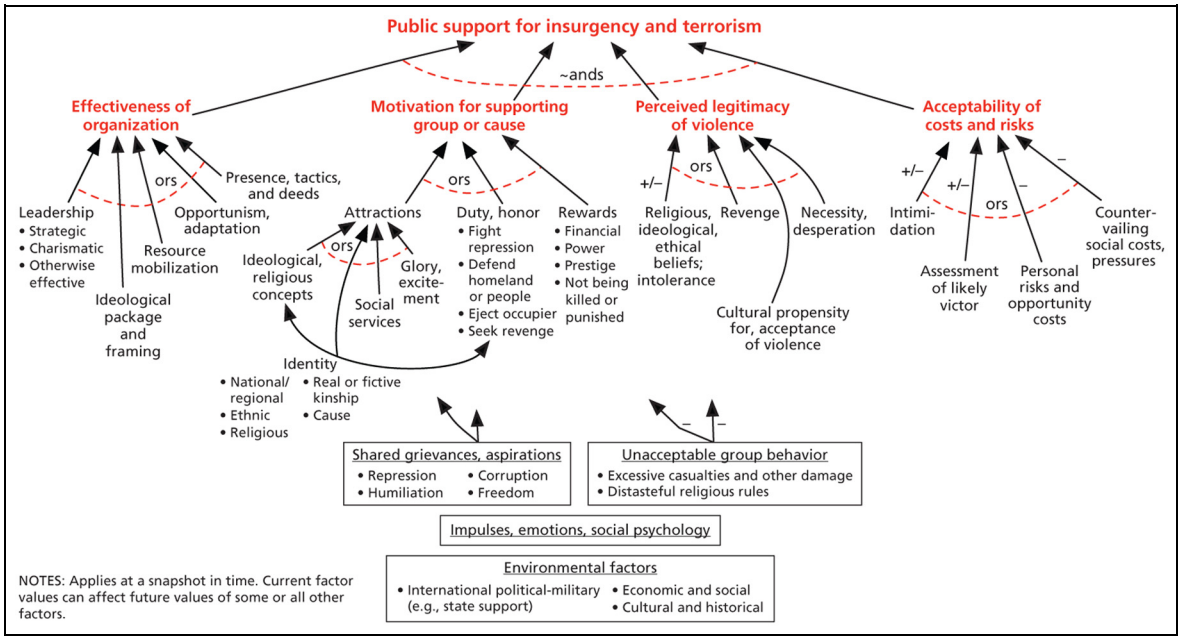

Figure 5. Factor-tree model that shows the relationships among factors that underlie the Public Support for Terrorism or Insurgency model of Davis and O'Mahony.[1]

Muslim world. They play up the perceived glory (*glry*) of supporting a cause that aims to redress these purported grievances. Religious beliefs and intolerance (*intl*) help increase the perceived legitimacy (*PLEG*) of violence against the West and against the many Muslims who do not share their Salafist views. Countervailing pressure (*scst*) discourages more support for al-Qa'ida. This countervailing pressure may occur in part because much of the public believe that al-Qa'ida is not likely to succeed and emerge as ultimate victors (*lvic*). This pressure in turn diminishes the acceptability of costs and risks (*ACR*) for al-Qa'ida activities. The parameters of this al-Qa'ida case study determined the relative causal edge weights in our FCM models.

3.1 The FCM-PSOT model

We cast the PSOT model as a FCM based on the original PSOT factor tree model. Then the directed edges of the factor tree became directed edges in the cognitive map. The signs of the factor tree links determined the signs of the FCM edges. The ambiguous factor-tree links (labeled "$+/-$" in Figure 5) defined weak bidirectional dependence between pairs of factors in the FCM. Factor tree nodes use "*~or*" or "*~and*" to aggregate their fan-in input signals. We used these different PSOT combination functions to specify analogous combination functions in the FCM. Factor tree nodes aggregate inputs by using functions from a predefined set of functions (Tables 2.3 and 2.4 in Davis and O'Mahony[1]). Nodes in a FCM apply nonlinear occurrence or activation functions to weighted linear combinations of their inputs as discussed above. The FCM-PSOT model used step functions (delayed or otherwise), logistic sigmoids, and the clamped linear activation functions.

The left panel of Figure 6 shows the direct FCM translation of the original PSOT model. The left panel of Figure 7 shows the intensity plot of the causal edge connection matrix for the FCM translation of the original PSOT model. The FCM-PSOT models in Figure 6 retain the general PSOT structure. But the edge weights are specific to the case of al-Qa'ida as researchers have reported in Figure 2.4 of Davis and O'Mahony[1] and Figure S.2 of Davis et al.[2]

3.2 The dynamic FCM-PSOT model

The PSOT factor tree gives a static snapshot of the state of public support for insurgency and terrorism. This makes the model useful for causal attribution at single points in time. This first-order static model does not require that we correctly identify causal cross-linkages among the factors. But the absence of cross-linkages can make long-run dynamic simulations misleading. An FCM with no feedback loops converges in at most L steps where L is the length of the longest chain in the FCM. Such feedforward or feedback-free models rarely give an accurate model of the real world and its causal interconnections. A dynamic

Table 1. Factors in the Public Support for Insurgency and Terrorism (PSOT) model.

Label	Full description
lead	Leadership, strategic or otherwise
pkg	Ideological package & framing
rsrc	Resource mobilization
opp	Opportunism & adaptation
pres	Presence, tactics, & deeds
EFF	Effectiveness of organization
reli	Ideological religious concepts
socs	Social services
glry	Glory, excitement
ATT	Attractions
duty	Duty & honor
rwrd	Rewards
MOTV	Motivation for supporting group, cause
intl	Religious, ideological, ethical beliefs; intolerance
rvng	Revenge
cprop	Cultural propensity for accepting violence
desp	Desperation, necessity
PLEG	Perceived legitimacy of violence
intm	Intimidation
lvic	Assessment of likely victor
prsk	Personal risk and opportunity cost
scst	Countervailing social costs & pressures
ACR	Acceptability of costs & risks
id	Identity
shgr	Shared grievances & aspirations
ugb	Unacceptable group behavior
env	Environmental factors
impl	Impulses, emotions, social psychology
hsucc	History of successes
mgtc	Management competence
prop	Propaganda, advertising
efdoc	Effectiveness of indoctrination, passing beliefs
hfail	History of failures
PSOT	Public support for insurgency and terrorism

time-varying PSOT model would need to identify such cross-linkages to track system behavior over time. A dynamic model gives causal attribution at snapshots in time *and* the ability to simulate long-run what-if scenarios.

Causal cross-linkages specify how sets of two or more factors co-vary in time. And they do so based on causal relationships. Causal cross-linkages are difficult to specify without insight into the causal laws that guide the related factors. Domain experts are the main source of these causal relations. The PSOT model relies on domain experts, extensive social science research, and validation to establish the snapshot relationships presented.[1,2,34] Specifying new causal cross-links in the model will require more such inputs from experts.

Our goal was to produce a *dynamic* causal model of public support for insurgency and terrorism in which new edges model covariation in time among model factors. The new edges needed grounding in subject-matter expertise. So we reviewed prior work on PSOT for information on factor covariation. We also consulted with PSOT authors and experts on the PSOT model for guidance on the new causal edges that we added. The new causal edges transformed the PSOT model from a static snapshot model into a dynamic simulation model. Figure 6 shows FCM versions of the old static and new dynamic PSOT model.

We now outline these changes to the original PSOT model. We first added a weak self-excitation feedback loop on the PSOT concept node because it is the highest-level concept node. This self-excitation loop modeled inertia in aggregate public opinion about insurgency and terrorism. This new feedback source induced a weak serial correlation in time in the PSOT concept node.

The next directed weak edges connected the top-level factors in Figure 5 from left to right: $EFF \rightarrow MOTV$, $MOTV \rightarrow PLEG$, and $PLEG \rightarrow ACR$. These directed causal edges made explicit an implicit point about O'Mahony and Davis's use of factor trees. Their factor-tree representation assumed a left-to-right dependence of the top-level factors that we have linked.[1,33] This implicit dependence made their factor tree more readable. The FCM model made this dependence explicit.

O'Mahony and Davis[1] discuss other dynamic augmentations to the PSOT model. They point to the following new factors. A history of successes or failures can affect motivation and perceived risks. We model this dependence with the two factors ''history of successes'' and ''history of failures.'' These two nodes exert opposing influence on *MOTV* and *prsk*. We split this history factor because traditional FCM models admit only positive values that represent the degree or intensity to which a concept occurs. And the effectiveness of the organization factor EFF partly determines the history of successes: $EFF \rightarrow hsucc$. Unacceptable group behavior *ugb* also influences motivation and effectiveness: $ugb \rightarrow MOTV$ and $ugb \rightarrow EFF$.

4 Simulation experiments and discussion

We first compare the behavior of the FCM-PSOT and dynamic FCM-PSOT models in the previous Sections 3.2 and 3.2. Then we examine methods for adaptively updating our fuzzy causal maps by using expert opinion or hard data or by using both. Factor tree models do not have access to these update methods.

4.1 Comparing the static and dynamic FCM-PSOT models

The PSOT FCM mimics the behavior of the original PSOT factor tree model because it is a direct FCM version of the factor tree model. This suggests that FCM models may be a richer classs of models than factor tree models because they act as feedback-laden supersets of tree or acyclic

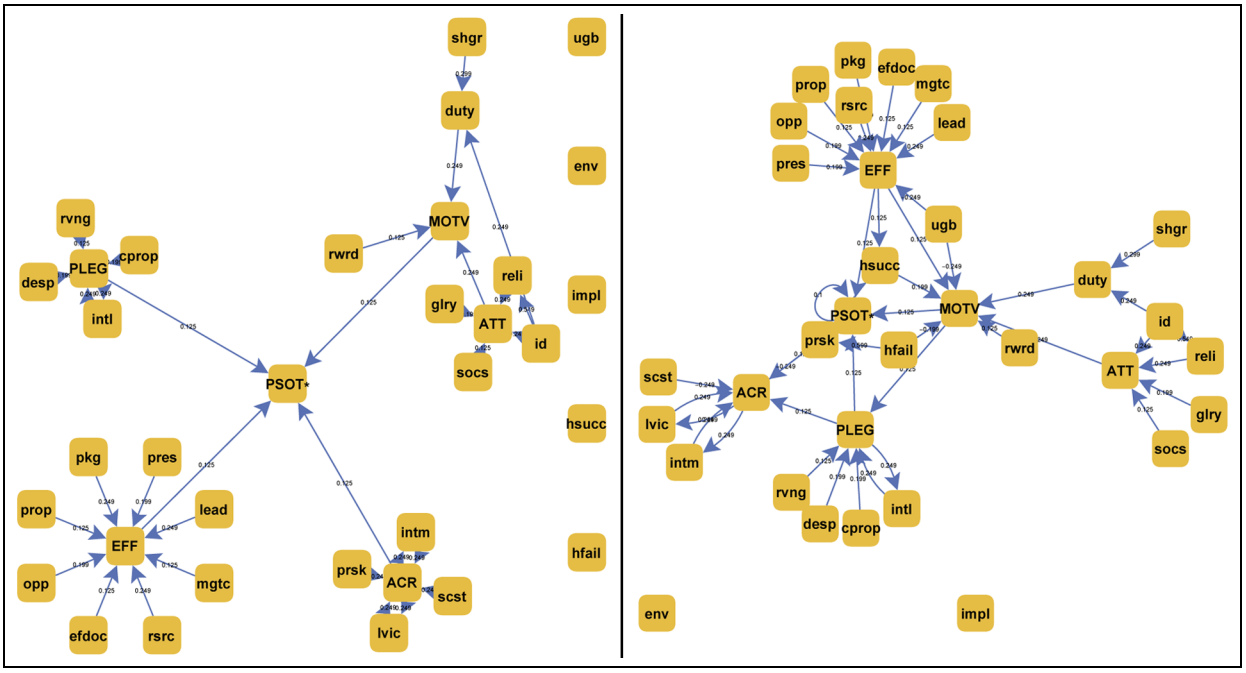

Figure 6. Fuzzy cognitive map implementations of the PSOT factor-tree model. Left: FCM digraph for the original static (acyclic) PSOT model. Right: FCM digraph of the dynamic PSOT model with cross-links. Table 1 gives the key for the concept-node labels in both FCMs. We based the new FCM edges in the right digraph on the findings in Davis and O'Mahony[1] or on expert input.

models. Further causal analyses may well find classes of factor trees that FCM models cannot easily capture.

Augmenting the PSOT with cross-links in the dynamic model allows richer representation of system dynamics. The dynamic and static FCM models agree on many inference tasks. Both models often converged to similar fixed-points or limit cycles given the same initial conditions. But examples of diverging behavior did appear. Network-analytic measures such as vertex degree and vertex centrality measures on the FCMs proved useful for inducing such divergent behavior.

Consider a hypothetical al-Qa'ida-like insurgent group that has similar PSOT model weights. Call the group the Salafist United (SU). SU's leadership is charismatic and competent as well as effective. Imagine Osama bin Laden or PKK's Abdullah Öcalan with careful ideological framing of their group's message and cause. This might be Salafism itself. We may assume also that people already inclined towards the group are culturally and ideologically comfortable with SU's violence. The group has embedded in a community that shares the group's strong Muslim identity. And SU's militant jihadi framing makes their cause attractive to many Muslims.

Suppose that SU has a history of failed operations despite its effective organization. Suppose further that the group has not made good use of political opportunities. The public believes that SU has a good chance of success. But SU routinely intimidates the public with violent or threatening tactics. These tactics impose social costs and personal risk on many members of the community. The group can bring only limited money and labor to bear on their campaign.

The scenario gives the following coding for the model. The factors *lead*, *hfail*, *pkg*, *pres*, *MOTV*, *cprop*, *intl*, *intm*, and *lvic* remain active throughout the evolution of this scenario. The factors rsrc, opp, prsk, and scst remain inactive during the simulation.

Both the static and dynamic FCM-PSOT models unfolded in time through the stated scenario constraints and initial conditions using logistic activation functions. And both models converged to fixed-point attractors instead of to limit cycles. The static model converged in 4 iterations to a fixed state that predicts little public support for SU. The dynamic FCM-PSOT converged in 11 iterations to a fixed state that predicted medium-to-high public support for SU. We also started the models from random initial states under the same constraints. Most of these perturbations died out before the FCMs converged to one of the fixed points.

The dynamic model predicts that a violent terrorist group can retain public support in a community that shares its religious or ideological beliefs and cultural propensities. It can retain that support in spite of a history of failure or a lack of resources. It can do so at a high cost or even

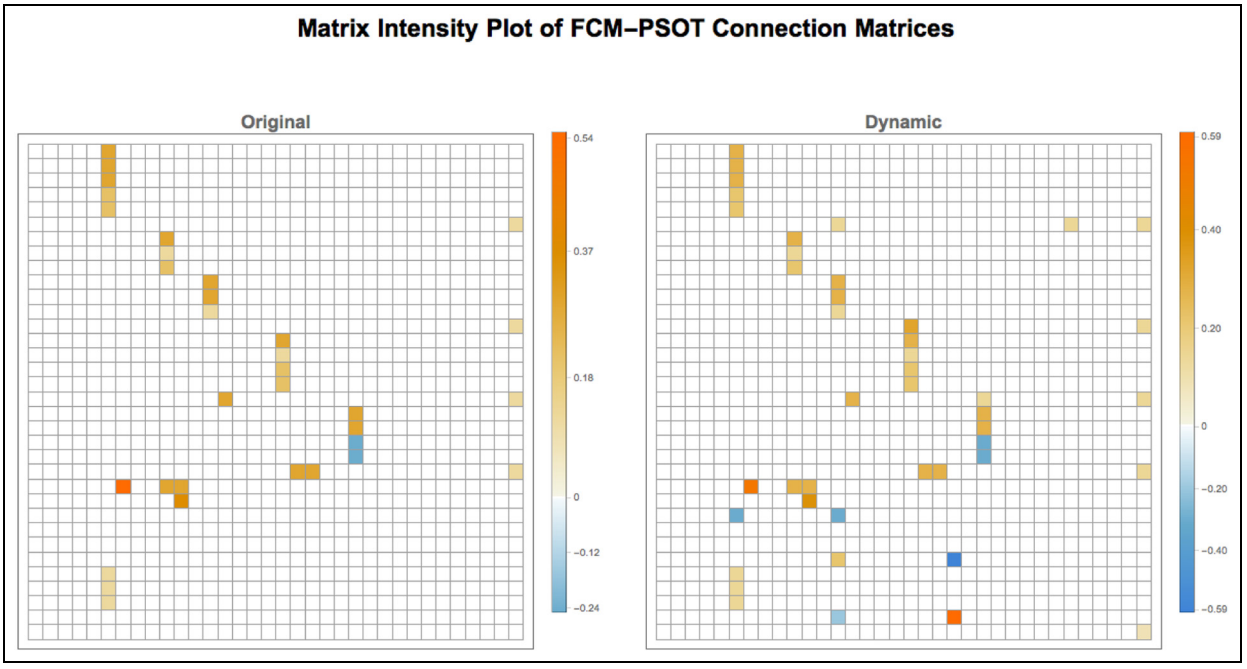

Figure 7. Causal-edge connection matrices **E** for the original static PSOT FCM (left) and the dynamic PSOT FCM (right). Each FCM's causal-edge or connection matrix **E** is the adjacency matrix for the FCM's fuzzy signed directed graph. Each square shows the fuzzy causal-edge value e_{ij} that denotes how much the ith concept C_i influences the jth concept C_j. The matrix entries e_{ij} in these FCMs are fuzzy values in the bipolar interval $[-1, 1]$. Blue squares represent negative causal influence and orange squares represent positive causal influence as the color bars indicate. White squares represent the absence of causal influence. These matrix intensity plots are larger-scale analogs of the matrix in Equation (22) except that they cover a larger set of concepts. The dynamic FCM-PSOT is marginally less sparse than the static FCM-PSOT because the dynamic FCM-PSOT asserts more directed causal-edge values between factors.

violence to the community in which the group acts. The static factor-tree-based model opposes this finding.

There is no easy way to validate the FCM-PSOT models because there is no clear ground-truth for such hypothetical simulations. But a good causal model can still give *pattern* predictions. It can help with exploratory analysis and detecting trends.

Consider how the dynamic FCM-PSOT model can help explore the effects of transient events. We omit the detailed vector-matrix operations and just qualitatively describe the resulting equilibrium FCM limit cycle.

Suppose the public does not support an insurgent group's cause ($MOTV = 0$). The group seems unlikely to succeed ($lvic = 0$). The public does not find the costs and risks of supporting this losing insurgency acceptable ($ACR = 0$).

Suppose now that some short-term event or shock occurs that both strongly motivates the public to support the insurgency's cause ($MOTV \to 1$) and that causes the public to believe that the group could win ($lvic \to 1$). Then the FCM-PSOT model predicts that the cost and risk of support will fall enough to become acceptable ($ACR = 1$). This leads the group to become bolder in its use of intimidation tactics ($intm \to 1$). But the group's gains will be short-lived and there will be no widespread public support for the insurgency if the shocking event is too short-lived to sustain public motivation for the cause ($MOTV \to 0$ and $lvic \to 0$). The public will once again find that the cost of support is too high. But suppose instead that the public motivation and the assessment of the likely victor remain steady ($MOTV = 1$ and $lvic = 1$). This could reflect an abusive government in rapid decline. Then the cost and risk will stay acceptable ($ACR = 1$). The public will eventually have sustained support for the insurgency ($PSOT \to 1$).

We can also argue for the value of the dynamic FCM model's findings by identifying insurgent organizations that resemble those in our scenario and that have managed to maintain public support. Resource mobilization problems pose a common concern for smaller extra-legal groups such as Uganda's LRA. Online radicalization appears to be replacing such resource mobilization problems. A related concern is a history of failure and the use of intimidation and violence against the locals as with Colombia's FARC rebels.

This argument by analogy is not itself a robust method of validation. It requires a comprehensive analysis of a larger sampling of terrorist and insurgent operations. The main way to build confidence in FCM inferences is to carefully build the causal knowledge model from representative evidence and then compare these what-if inferences or predictions with observed patterns or outcomes.

4.2 Causal knowledge updating by averaging expert responses

The tree-structure of the original PSOT model implies that it inherits a key structural limitation of AI decision trees. There is no easy or natural way to combine or fuse several expert trees into a representative knowledge structure that is still a tree. Cycles too easily appear in general as the number or sample size m of fused experts increases. Combining even a small number of expert trees is likely to produce some cycles and hence feedback loops in the combined knowledge structure.

Some form of ad hoc cycle clipping must ensure that the combined trees produce a tree. But removing causal cycles removes some of the very expert knowledge that the tree structure tries to capture. And it does so solely to maintain the tree structure. FCM models are in this sense at least as expressive as factor tree models. They also benefit from the strong law of large numbers if the combined experts behave at least approximately as independent and identically distributed knowledge sources.[12,14]

Figure 4 gives a simple demonstration of how such model updates can occur. The mixture or convex combination of FCMs creates a new fused FCM as the weighted average of the FCMs' augmented simply matrices. Users can add new factors at will. Each new factor converts the n-by-n adjacency matrices into $n+1$-by-$n+1$ adjacency matrices. This amounts to adding a new zero-padded row and column to an adjacency matrix if its corresponding FCM does not include the factor as a concept node.

This fusion-averaging technique may not directly account for such effects as active sabotage or extreme variance in expert opinions. Highly variable expert inputs will tend to produce a highly variable FCM causal knowledge base. There may be no benefit from combining expert edge values that approximate thick-tailed probability densities. Cauchy probability bell curve closely resemble normal probability bell curves. Cauchy bell curves have slightly thicker tails that give rise to far more variable realizations. But the sample average of Cauchy random variables is itself a Cauchy random variable. So there is no benefit or decrease in system variance whatsoever in this thick-tailed case. The combined result has the same infinite variance that any one of the individual Cauchy edges has. Combining knowledge sources with even thicker-tailed probability densities can produce variability even more extreme than the variability of any of the combined knowledge sources.

Model averaging for FCMs helps update our knowledge of causal edge values. But we may also want to update our knowledge of the causal concepts $C_k(t)$ as well. We can use data from expert opinion surveys, from direct time-series data on measurable factors, or from indirect instrumental variables linked to the factors of interest. Such instrumental variables include social media trends, Google trends, and topic modeling on news corpuses.

4.3 Learning causal structure with differential Hebbian learning

We draw a learning distinction between factor correlation and factor covariation. Consider first factor correlation.

Directed causal edges induce correlations between linked factors. These correlations themselves need not indicate a causal dependence. We can estimate these correlations from observations given enough samples and effort. This estimation of correlative links requires time-series data about variation in the factors. The two learning approaches below assume for simplicity that there are no time delays between factors. Users can easily insert such time delays as needed.

We can learn causal edge strengths through the *concomitant activation* among the set of factor pairs. This approach assumes that events (factor activities) are more likely to involve a causal connection if the events occur together.[13,37,38] This suggests the well-known Hebbian correlation learning law (neurons that fire together wire together) for training neural network synaptic weights[13]:

$$\dot{e}_{ij} = -e_{ij} + C_i C_j \tag{26}$$

where $\dot{x}$ denotes the time derivative of the signal x. The passive decay term $-e_{ij}$ stabilizes the learning in the differential-equation model. It also models a ''forgetting'' constraint that helps the network prune inactive connections. The product term $C_i C_j$ directly models concomitant correlation.

We can alternatively use *concomitant variation*[39] in time between factors as partial evidence of a causal relation between those factors. Suppose that the data indicate that increases in factor C_i occur at the same time as increases in the factor C_j. This concomitant increase suggests that the edge value e_{ij} should be positive. Suppose similarly that decreases in C_i occur with decreases in C_j. Then such concomitant decrease suggests a negative causal-edge value e_{ij} (even a slight time lag between the two concept nodes can indicate the direction of causality in practice). Such concomitant variation leads to the *differential* Hebbian learning law[12,13,38]:

$$\dot{e}_{ij} = -e_{ij} + \dot{C}_i \dot{C}_j \tag{27}$$

Both Hebbian learning and Differential Hebbian Learning (DHL) methods can learn causal edge values in a FCM. But Hebbian learning grows spurious causal relations between any two concept nodes that occur at the same time. This quickly leads to a matrix of nearly all unity values if most of the nodes are active. DHL correlates node velocities and thus has a type of arrow of time built into it. DHL correlates the signs of the time derivatives. So it grows a positive causal edge value e_{ij} if and only if the concept nodes C_i and C_j both increase or both decrease. It grows a negative edge value if and only if one of the nodes increases and the other decreases.

We can combine both learning laws to give a more general version of DHL[12]:

$$\dot{e}_{ij} = -e_{ij} + C_i C_j + \dot{C}_i \dot{C}_j \tag{28}$$

This hybrid learning law fills in expected values for edge weights when there is no signal variation in the factor set.[40] It takes advantage of the relatively rarer variation events to update the edge weights. It tends to produce limit cycles or other attractors but can produce fixed-point attractors given some strong assumptions.[12,13]

We can use the DHL learning scheme to infer causal weights in FCM-PSOT models if we have access to adequate time-series data. Such data can again come from expert opinion surveys, from direct time-series data on measurable factors, or from indirect instrumental variables linked to the factors of interest: social media trends, Google trends, or topic modeling on news corpuses. Figure 3 shows a DHL learning path for a single causal edge value. The algorithm used data under a hypothetical relationship for the uncertain link $lvic \rightarrow ACR$ in the PSOT model. Complete concept-node data allows the DHL algorithm to learn the causal edge matrix $\mathbf{E}$. We found that DHL training gave a close approximation of the true causal-edge values after only a few iterations.

Our simulations with adaptive FCMs used the following discretized version of the DHL learning law[10] in Equation (27):

$$e_{ij}(t+1) = \begin{cases} e_{ij}(t) + \mu\left[\Delta C_i(t)\Delta C_j(t) - e_{ij}(t)\right] & \text{if } \Delta C_i(t) \neq 0 \\ e_{ij}(t) & \text{else} \end{cases} \tag{29}$$

where $\Delta C_k(t) = C_k(t) - C_k(t-1)$.

We note that we can fuse soft and hard knowledge sources through the above averaging technique in Equation (25). Let $\mathcal{E}_{\text{data}}$ denote the data-driven FCM. Let $\mathbf{E}_{\text{exp}}$ denote the expert-elicited FCM. Then the fused causal-edge matrix $\mathbf{E}_{\text{fusion}}$ is a simple mixture of the two edge matrices:

$$\mathbf{E}_{\text{fusion}} = \omega_{\text{data}}\mathbf{E}_{\text{data}} + \omega_{\text{exp}}\mathbf{E}_{\text{exp}} \tag{30}$$

Then Equation (29) or some other statistical learning law can continue the adaptation process using new numerical data or occasional opinion updates from experts.

5 Conclusions

We developed static and dynamic FCM versions of a factor tree model of public support for insurgency and terrorism. The FCM models allow forward-chaining causal inference as well as updates based on numerical training data or expert opinions. Their underlying matrix structure permits natural knowledge fusion that tends to improve with the number of combined experts.

Current FCM learning techniques involve two major limitations. FCMs do not easily permit backward chaining to answer which input cause produced an observed output effect. Users cannot simply run the FCM in reverse because of the node nonlinearities. The best that current techniques allow is to find one of the many input states that map to an observed output equilibrium. Future research needs to address this limitation with new inferencing or other techniques. The second limitation is even more challenging. Current adaptive techniques infer and tune causal-edge values only for a known set of concept nodes. An open research question is how to use data-based techniques to infer new or missing concept nodes in large-scale FCM causal models.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

References

1. Davis PK and O'Mahony A. A computational model of public support for insurgency and terrorism: A prototype for more-general social-science modeling. Technical report, RAND Corporation, 2013.
2. Davis PK, Larson EV, Haldeman Z, et al. *Understanding and influencing public support for insurgency and terrorism*. Santa Monica, CA: RAND Corporation, 2012.
3. Ibrahim R. *The Al Qaeda reader: The essential texts of Osama Bin Laden's terrorist organization*. New York: Broadway Books, 2007.
4. Nawaz M. *Radical: My journey from Islamist extremism to a democratic awakening*. New York: Random House, 2015.
5. Glykas M (ed.) *Fuzzy cognitive maps: Advances in theory, methodologies, tools and applications*. Berlin: Springer, 2010.
6. Papageorgiou E. *Fuzzy cognitive maps for applied sciences and engineering: from fundamentals to extensions and learning algorithms*. Berlin: Springer, 2013.

7. Kosko B. Fuzzy cognitive maps. *Int J Man Mach Stud* 1986; 24: 65–75.
8. Dickerson JA and Kosko B. Virtual worlds as fuzzy cognitive maps. *Presence* 1994; 3: 173–189.
9. Kosko B. Fuzzy systems as universal approximators. *IEEE Trans Comput* 1994; 43: 1329–1333.
10. Kosko B. *Fuzzy engineering*. Upper Saddle River, NJ: Prentice Hall, 1996.
11. Osoba O, Mitaim S and Kosko B. Bayesian inference with adaptive fuzzy priors and likelihoods. *IEEE Trans Syst Man, Cybern, Part B Cybern* 2011; 41: 1183–1197.
12. Kosko B. Hidden patterns in combined and adaptive knowledge networks. *Int J Approximate Reasoning* 1988; 2: 377–393.
13. Kosko B. *Neural networks and fuzzy systems: A dynamical systems approach to machine intelligence*. Upper Saddle River, NJ: Prentice Hall, 1991.
14. Taber R, Yager RR and Helgason CM. Quantization effects on the equilibrium behavior of combined fuzzy cognitive maps. *Int J Intell Syst* 2007; 22: 181–202.
15. Davis PK, Manheim D, Perry WL, et al. *Causal models and exploratory analysis in heterogeneous information fusion for detecting potential terrorists*. Santa Monica, CA: RAND Corporation, 2015.
16. Davis PK, Manheim D, Perry WL, et al. Using causal models in heterogeneous information fusion to detect terrorists. In: *2015 winter simulation conference (WSC)* (eds L Yilmaz, WKV Chan, I Man, et al.), Huntington Beach, CA, 6–9 December 2015, pp.2586–2597. Piscataway, NJ: IEEE.
17. Kosko B. Fuzzy knowledge combination. *Int J Intell Syst* 1986; 1: 293–320.
18. Taber R. Knowledge processing with fuzzy cognitive maps. *Expert Syst Appl* 1991; 2: 83–87.
19. Zadeh LA. Fuzzy sets. *Inform Control* 1965; 8: 338–353.
20. Zadeh LA. Outline of a new approach to the analysis of complex systems and decision analysis. *IEEE Trans Syst Man Cybern* 1973; 3: 28–44.
21. Kosko B and Isaka S. Fuzzy logic. *Sci Am* 1993; 269: 62–7.
22. Zadeh LA. Probability measures of fuzzy events. *J Math Anal Appl* 1968; 23: 421–427.
23. Kissinger HA. Starting out in the direction of Middle East peace. *Los Angeles Times*, Summer 1982.
24. Osoba O and Kosko B. The noisy expectation-maximization algorithm for multiplicative noise injection. *Fluctuation Noise Lett* 2016; 15: 1650007.
25. Kosko B. Fuzzy entropy and conditioning. *Inform Sciences* 1986; 40: 165–174.
26. Kosko B. Probable equality, superpower sets, and superconditionals. *Int J Intell Syst* 2004; 19: 1151–1171.
27. Lewis D. Probabilities of conditionals and conditional probabilities. *Philos Rev* 1976; 85: 297–315.
28. Lewis D. Probabilities of conditionals and conditional probabilities – part two. *Philos Rev* 1986; 95: 581–589.
29. Kosko B. Bidirectional associative memories. *IEEE Trans Syst Man Cybern* 1988; 18: 49–60.
30. Predd JB, Osherson DN, Kulkarni SR, et al. Aggregating probabilistic forecasts from incoherent and abstaining experts. *Decis Anal* 2008; 5: 177–189.
31. Pearl J. *Causality: Models, reasoning, and inference*. New York: Cambridge University Press, 2009.
32. Cooper GF. The computational complexity of probabilistic inference using Bayesian belief networks. *Arti Intell* 1990; 42: 393–405.
33. Davis PK. Primer for building factor trees to represent social-science knowledge. In: *Proceedings of the winter simulation conference* Phoenix, AZ, 11–14 December 2011, pp.3121–3135. Piscataway, NJ: IEEE.
34. Davis PK, Cragin K, Noricks D, et al. *Social science for counterterrorism: Putting the pieces together*. Santa Monica, CA: RAND Corporation, 2009.
35. Snow DA, Soule SA and Kriesi H. *The Blackwell companion to social movements*. Malden, MA: John Wiley & Sons, 2008.
36. Davis PK and Bigelow JH. *Experiments in multiresolution modeling (MRM)*. Santa Monica, CA: RAND Corporation 1998.
37. Hebb DO. *The organization of behavior: A neuropsychological approach*. New York: John Wiley & Sons, 1949.
38. Kosko B. Differential Hebbian learning. *AIP Conf Proc* 1986; 151: 277–282.
39. Mill JS. *A system of logic, ratiocinative and inductive: Being a connected view of the principles of evidence and the methods of scientific investigation*. London: John W. Parker, 1843.
40. Kosko B. Unsupervised learning in noise. *IEEE Trans Neural Networks* 1990; 1: 44–57.

Author biographies

Dr. Osonde A. Osoba is a researcher at the RAND Corporation and a professor at the Pardee RAND Graduate School, Santa Monica, CA, USA.

Dr. Bart Kosko is a professor of Electrical Engineering and Law at the University of Southern California, Los Angeles, CA, USA.