Introduction and Historical Background

Global crises have a significant impact on systems across the world including financial institutions, supply chains, travel, etc. These crises often reveal critical dependencies, protocols, and points of failure within and across global systems. The after-math of the 2008 financial crisis, for example, revealed a precarious dependency on subprime mortgages and rating agencies. The 1973 OPEC crisis revealed dependencies on U.S oil imports from OPEC nations. Similarly Covid-19 has uncovered countless key dependencies on personal protective equipment (PPE), emergency care resources, etc. The ability to prematurely model these risks/volatilities and identify interdependence effectively can help enact proactive policies rather than the historical trend of deploying reactive responses.

The first step in approaching this problem is to characterize global systems. Crises like Covid-19 can be defined by complex adaptive systems, where the whole cannot be described simply by the sum of the parts. As John H. Miller succinctly puts it, "One and one may well make two, but to really understand two we must know both about the nature of 'one' and the meaning of 'and.'" (Miller 2007) There are two primary motivators for this approach. First, global systemic risk modeled as a CAS is cognizant of emergent features allowing for identification of behaviors that would have otherwise been difficult to observe. Second, CAS's are composed of smaller individual agents which allows for larger systems (or intersection of systems) to be subdivided into their smaller defining components.

Fuzzy Cognitive Maps

Background and Implementation

CASs can effectively be modeled by Fuzzy Cognitive Maps (FCMs)(Tlili, Chikhi, 2012). Fuzzy Cognitive MFCMs combine fuzzy logic with the structure of directed graphs. Fuzzy logic is a subset of fuzzy mathematics where a proposition (e.g Akash is old) is resolved with a non-binary degree of "correctness" (typically any real values between 0 and 1 inclusive). FCM's are represented by a set of nodes and edges where pairs of nodes are connected by unidirectional edges. Nodes represent various concepts (e.g population, law enforcement, etc) and have an internal fuzzy value representing a degree of the concept. Edges are weighted connections that indicate causality between nodes. An edge with a positive weight value indicates a positive causation between the two connected nodes, whereas a negative weight indicates a negative causation. This can be seen through Figure 1.

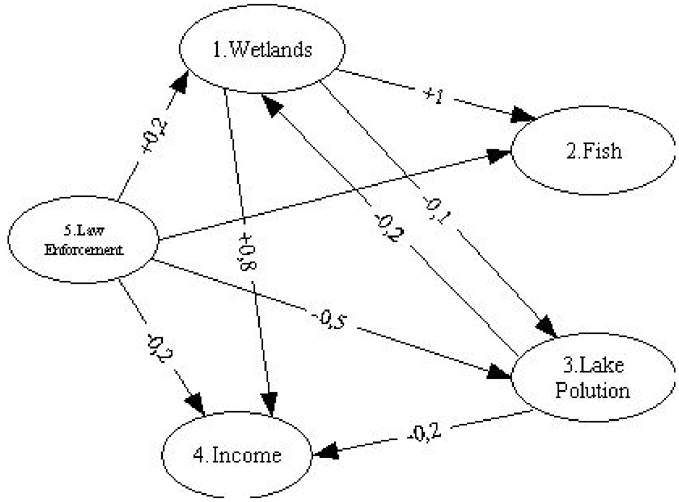

Figure 1: Özesmi & Özesmi 2004 FCM where 5 concepts (wetlands, fish, lake pollution, income, and law enforcement) are connected by unidirectional weighted connections.

Fuzzy Cognitive Maps with defined weights as shown in Figure 1 can be used to forecast future scenarios. This can be accomplished by creating a vector, A^t, representing the initial state of the concepts at time t. Simulating the progression from t to t+1 can be achieved by updating the value of each concept using the formula shown in Figure 2.

$$A_i^{(t+1)} = f(\sum_{j=1}^{N} A_j^{(t)} w_{ji})$$

Figure 2: Source - Learning rule (Gregor, Groumpos, 2013)

w_{ji} is the scalar weight pointing from concept j to concept i. N is the total number of concepts in the graph. The subscript for the vector A^t and $A^{(t+1)}$ indicate the specific node being referenced (a numerical index for concepts). Function f is a activation function (also called a squashing function) that fits the dot product taken with $\sum_{j=1}^{N} A_j^{(t)} w_{ji}$ into a specific range (typically [0,1] or [-1,1]). Common activation functions include: sigmoid ($f(x) = \frac{1}{1+e^{-x}}$), hyperbolic tangent or tanh ($f(x) = tanh(x)$), and rectified linear unit or ReLU ($f(x) = max(0, x)$). The activation function can be chosen empirically based on performance or the context of the model.

	Wetlands	Fish	Lake Pollution	Income	Law Enforcement
Wetlands	0.0	+1	-0.1	+0.8	0.0

Fish	0.0	0.0	0.0	0.0	0.0
Lake Pollution	-0.2	0.0	0.0	-0.2	0.0
Income	0.00.0	0.0	0.0	0.0	0.0
Law Enforcement	0.2	0.5	-0.5	-0.2	0.0

Figure 3: Representing the FCM from Figure 1 in weight matrix form

Defining Weight Matrices

For an FCM to be used in forecasting, a weight matrix (e.g Figure 3) describing the causality between concepts must be determined. This can be realized through two methods: construction of the matrix through domain experts or automated learning algorithms. Using experts in a certain domain allows for past research and proven casualties to be represented in the map effectively. Tools like MentalModeler (http://www.mentalmodeler.org/) allow researchers to use a graphical interface to design semi-quantitative FCMs. Users define the relevant concepts, outline a weight matrix, and specify their confidence in each weight. This approach is particularly useful when there is limited data in the field or casualties simply cannot be determined through quantitative data. However, expert defined matrices make it difficult to scale the number of concepts in an FCM due to the high dependency on manual data analysis. Further, the accuracy of each model is limited by the casualties defined by experts and is vulnerable to various biases.

Learning algorithms eliminate the need for expert insight by tuning weight matrices using data. A single sample of training data consists of two vectors: an initial and final state for each concept in an FCM. Time series data for each concept in a map can be used and reformatted into pairs of initial and final state vectors. There are a handful of training algorithms available that typically fall into two categories: supervised or unsupervised. FCM research over the last 20 years has had a greater focus on unsupervised learning methods like Hebbian learning (based loosely on the idea that neurons that fire together wire together) and genetic algorithms. More uncommon FCM training methods use supervised learning with gradient descent. (See appendix for more details on Hebbian learning and gradient descent methods). This approach eliminates expert bias from a weight matrix and can scale the number of concepts more effectively. However working with learning algorithms requires substantial data to generate accurate weight values. This scale and granularity of data can be difficult to acquire for larger global systems (other than financial data which can be used as a proxy for various concepts).

Usage/Utility of an FCM

Once an FCM is defined with an accurate weight matrix there are countless useful insights that can be derived.

- If a learning algorithm is used, the weight matrix can be examined to identify casualties.
- Centrality of nodes can be calculated (e.g degree, closeness, betweenness) to determine how integral they are in the function of the entire graph. This can be a useful indicator in identifying strong interdependency and points of failure, as failure of highly central nodes can cause cascading changes/failures across the entire map.

  - Scenarios can be forecasted using the equation in Figure 2. A hypothetical scenario can be used as an initial state vector and simulated to predict the output vector at time t+1. A perturbation in certain concept values can be simulated to see the impact they have on the rest of the system. This can be used as an effective measure of volatility/stability of an FCM.

- Historical failure in systems reveals points of failure and interdependence
  -
- Complex Adaptive Systems
  - Emergent features
    - Use examples from cases above
  - Modeling it as a directed graph of causalities
- Cognitive Maps
  - Intro - Concepts/Connections
  - 'Fuzzy' cognitive maps
    - Expert designed
    - Computationally generated
    - Evaluation metrics+usage
- Other approaches and considerations
  - RNN/LSTM

Introduction - 1 pages

Goals/Problem Statement - 1 page

Background - 5 pages

- LSTM - 1 pages
- FCM - 2 pages
- ??? - 2 pages

Procedure - 5 pages

- Evaluation Metrics - 1 page
- Datasets + Processing - 1 page
- Algorithms - 3 pages

Evaluation Metrics - 1 page

Results - 2 pages

Conclusion - 1 page