To appear: *World Conference of the International Fuzzy Systems Association* (*IFSA 2023*) -- *August 2023*

Searching for Causal *Phantom Nodes* in Fuzzy Cognitive Maps

Akash Kumar Panda
Signal and Image Processing Institute
Department of Electrical and Computer Engineering
University of Southern California
Los Angeles, California

Bart Kosko
Signal and Image Processing Institute
Department of Electrical and Computer Engineering
University of Southern California
Los Angeles, California
kosko@usc.edu

***Abstract*—We call causal variables *phantom nodes* in a fuzzy cognitive map (FCM) if they affect the FCM but the FCM does not include them in its web of nodes and causal edges. Supervised and unsupervised learning schemes can estimate the causal connections to a phantom node based on how the learned causal edges affect the FCM's equilibrium attractors. We illustrate this technique with gradient descent on the well-studied dolphin FCM. The process starts with the 5-node FCM and finds its equilibrium limit cycles. Then we remove one of the nodes and treat it as a phantom node and train. Gradient learning trains the augmented FCM with samples from the FCM edges using the known target limit cycles. The simulations used squared error as the performance measure but other measures can apply. The causal learning methodology extends to more than one phantom node but at greater computational cost.**

***Index Terms*—causal inference, fuzzy cognitive maps, hidden variables, causal learning, AI hallucinations**

I. The Search for Phantom Nodes

How do we detect and model missing causal variables in a causal model?

We call such missing or hidden variables *phantom nodes* in rough analogy with the neural hallucination of missing phantom limbs [1]. We here search for phantom nodes in feedback fuzzy cognitive maps (FCMs) [2], [3].

Detecting phantom nodes involves a type of causal or AI hallucination [4] in a process that resembles adding noise to a nonlinear system in stochastic resonance [5]. A small amount of causal hallucination can help estimate a missing relevant node while too much hallucination tends to obscure it.

Even detecting causal phantom nodes is far from trivial. These hidden or phantom variables can causally affect the other variables in a causal model. This can change the causal model's dynamics and thus change its equilibrium causal predictions [6]–[8].

These changes can change the dynamics of a feedback FCM and thereby change its equilibrium attractors such as its fixed points or limit cycles. These attractors act as the FCM's what-if predictions given input stimuli. This change in dynamics and equilibrium structure can hint at the presence and nature of the phantom variables. That is the idea behind the limit-cycle-based algorithm in this paper that estimates the causal structure of a phantom node from known target limit cycles. The known limit cycles allow validation of the predicted causal edges to and from the phantom concept node. They may correspond in practice to a user's desired policy outcome.

Fuzzy Cognitive Maps (FCMs) model causality in complex feedback dynamical systems [2], [3], [3], [9]–[15]. The causal variables of the FCM often come from domain experts who model system dynamics. The matrix structure of their underlying signed digraph structure allows users to combine or mix any number of related FCMs (augmenting with appropriate rows and columns of zeroes for missing nodes) into a fused or combined FCM. These combined FCMs can have much richer dynamics and equilibria than do the individual FCMs [11], [15].

We can search for these missing or phantom causal variables by adding them to the FCM causal system and then learning from the dynamics of the original system. This paper illustrates the technique with supervised gradient descent on the squared error between the current limit cycles and the known target limit cycles. Figure 1 illustrates this approximation technique and explains the limit-cycle approach that we use to estimate a phantom node. Time-series data can also train differential Hebbian learning [9], [16] or other unsupervised or supervised causal learning algorithms.

Section II describes the FCM feedback model used in this paper. The FCM models complex causal relationships among expert-defined variables in a dynamical system as a directed graph with fuzzy causal edges. The nodes of the graph define the causal variables. The directed edges define the partial signed causality between the variables. The FCM itself is a cyclic signed fuzzy directed graph. So it is a feedback dynamical system whose equilibrium attractors can be fixed points, limit cycles, or even aperiodic or chaotic attractors for maps with sufficiently complex node nonlinearities and edge structure. Section II also introduces the well-studied Dolphin FCM and its behavior as an example [10], [16].

Section III describes the concept of phantom nodes and also describes the FCM augmented with phantom nodes. This section also gives an example of a phantom node in the Dolphin FCM.

Section IV describes a gradient-based learning scheme to estimate phantom nodes. The phantom-node-augmented FCM learns from the dynamics of the system it models and tries to match its limit cycles. Section V presents the experimental

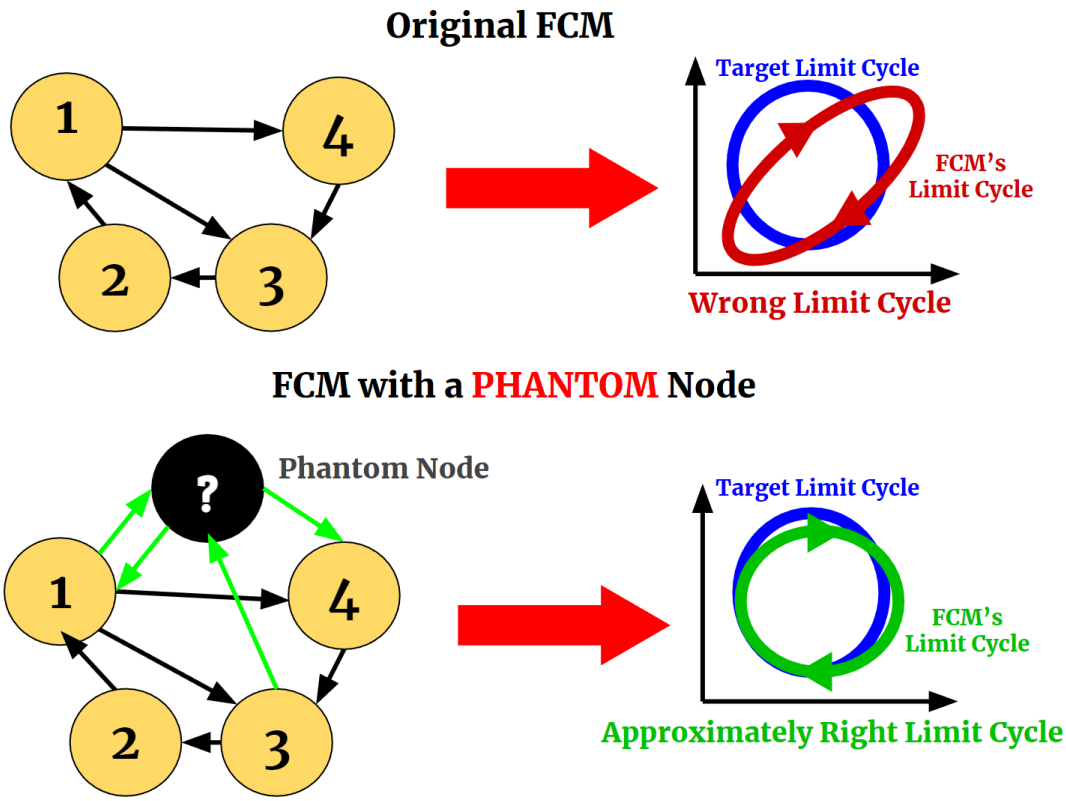

Fig. 1. Causal phantom nodes augment a fuzzy cognitive map (FCM) and help approximate the FCM dynamics and its target equilibrium limit cycles. The figures on the left show the FCMs and the figures on the right show their corresponding limit cycle. The figures on the top correspond to the original FCM with 4 nodes. The figures on the bottom correspond to the augmented FCM with 4 observable nodes and one phantom node. The original FCM does not approximate the limit cycles of the system it models. The FCM with the phantom node can approximate the target limit cycles with its own limit cycles.

results. The augmented FCM did approximate the limit cycles despite not approximating the phantom edges.

Figure 2 shows a simple version of the known Dolphin FCM that models how a pod of dolphins responds to the presence of a survival threat such as a shark. Figure 3 shows the different nonlinear activation functions that the FCMs can use. We use a rectified hyperbolic tangent (ReTanh) function as our activation function because it is continuous and also ensures that the node remains completely inactive unless it gets a positive stimulus. Figure 4 shows a few limit cycles of the dolphin FCM. These limit cycles show that the Dolphins will alternate between periods of rest and running away from the survival threat. Figure 5 shows the change in the limit cycles that correspond to the change in the steepness of the activation function.

Figure 6 shows the Dolphin FCM with one phantom node. Figure 7 shows the effect of a phantom node "SURVIVAL THREAT" on the dolphin FCM's limit cycles. Figure 8 shows the effect of detecting "bad" phantom nodes. Figure 9 shows the gradient-based phantom-node causal learning process as a flowchart. Figure 10 shows a limit cycle of the Dolphin FCM and compares it to the approximated limit cycles of the phantom-node augmented FCM. We present the results for just one of the Dolphin-FCM phantom nodes since other choices gave similar results.

II. Fuzzy Cognitive Maps (FCMs)

A Fuzzy Cognitive Map (FCM) connects concept nodes C_i through directed fuzzy or partial causal edges e_{ij}. FCMs allow feedback and thus nontrivial equilibrium dynamics unlike the feedforward acylic DAGs found in Bayesian belief trees [17].

The edge weight $e_{ij} \in [-1, 1]$ gives the degree to which C_i causes C_j in how much it increases or decreases C_j:

$$e_{ij} = \text{Degree}(C_i \to C_j) \tag{1}$$

The FCM is a dynamical system and so the node values $C_i \in [0, 1]$ change with time. Consider a discrete-time FCM. The $C_j(t+1)$ depends on the values $C_i(t)$ of nodes connected to C_i. It also depends on the edges e_{ij} that connect C_i to C_j:

$$C_j(t+1) = \Phi\left(\sum_{i=1}^{m} C_i(t) e_{ij}\right) \tag{2}$$

where n is the number of nodes and Φ is a nonlinear function that maps to $[0, 1]$. This nonlinear activation function Φ may be a threshold function, a logistic sigmoid function, a rectified hyperbolic tangent (ReTanh) function, or some other nonlinear function. We used the ReTanh function:

$$\Phi(x) = \max(0, \tanh(cx)) \tag{3}$$

where c controls the steepness of the activation function Φ.

Consider the 5-node FCM in Figure 2. The FCM describes the group behavior of a dolphin pod in the presence of a survival threat such as a shark. Then the causal-edge matrix E can represent this dolphin FCM:

$$E = \begin{pmatrix} 0 & 1 & 0 & -1 & 0 \\ 0 & 0 & 1 & 0 & -1 \\ 0 & -1 & 0 & 1 & -1 \\ 1 & 0 & -1 & 0 & 1 \\ -1 & 1 & 0 & -1 & 0 \end{pmatrix} \tag{4}$$

Suppose a shark appears. This corresponds to $C_4(0) = 1$. Let all other nodes be 0. Then the vector $\begin{pmatrix} 0 & 0 & 0 & 1 & 0 \end{pmatrix}$ is the initial state of the FCM. Equation (2) updates the FCM to the state $\begin{pmatrix} 1 & 0 & 0 & 0 & 1 \end{pmatrix}$ using a simple threshold at zero. The FCM keeps updating and enters the limit cycle $\begin{pmatrix} 0 & 0 & 0 & 1 & 0 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & 0 & 0 & 0 & 1 \end{pmatrix} \rightarrow \begin{pmatrix} 0 & 1 & 0 & 0 & 0 \end{pmatrix} \rightarrow \begin{pmatrix} 0 & 0 & 1 & 0 & 0 \end{pmatrix} \rightarrow \begin{pmatrix} 0 & 0 & 0 & 1 & 0 \end{pmatrix}$. Figure 4 shows a some examples of this FCM's limit cycles.

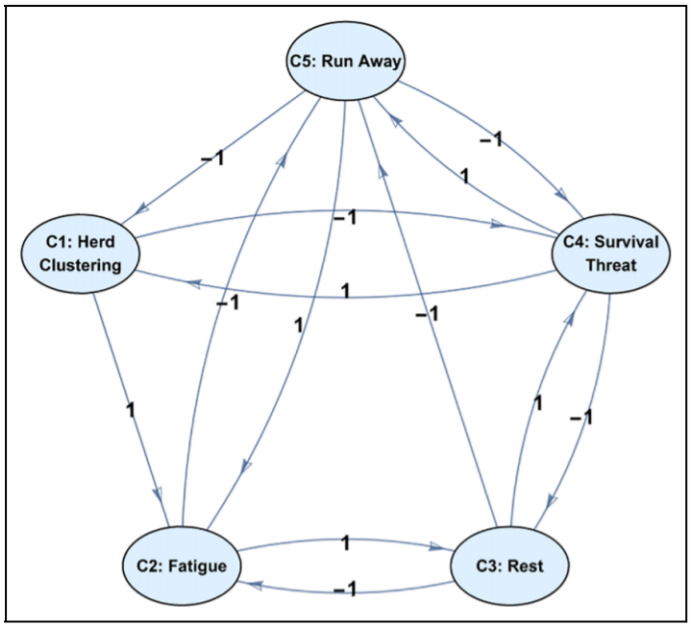

Fig. 2. Dolphin FCM of dolphins near a threat such as a shark.

We use the truncated or Rectified hyperbolic Tangent (Re-Tanh) function Φ to ensure that $\Phi(x) = 0$ when $x \leq 0$. The dolphin FCM with rectified bipolar-logistic functions gives the same limit cycles as the threshold-FCM. The limit cycles fade out when c is small. Figure 5 shows the time evolution of this FCM for initial state $\begin{pmatrix} 0 & 0 & 0 & 1 & 0 \end{pmatrix}$ and for different values of c. It also compares these limit cycles to the limit cycles of the threshold-FCM.

III. Causal Phantom Nodes in FCMs

Consider a system where the experts do not know or cannot estimate all the causal links. Then their FCM (likely the result of a weighted mixing of their edge matrices) may well not match the system's true dynamics. So there may be one or more hidden or phantom concept nodes that the experts missed and that causally connect to the observable nodes and that may affect their dynamics. These phantom edges augment the edge

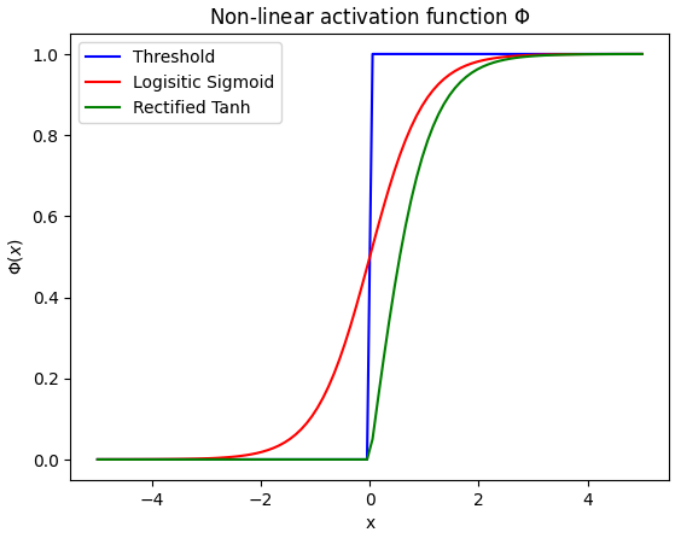

Fig. 3. The choices for the non-linear function. The simple threshold is in blue. The logistic sigmoid is in red. The rectified hyperbolic tangent Φ is in green. Here $c = 2$.

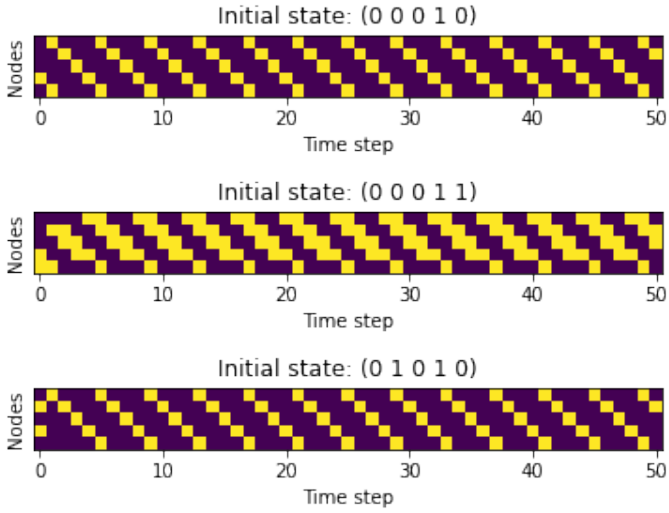

Fig. 4. Limit cycles from the Dolphin FCM. The time step is along the x-axis. The nodes are along the y-axis. The images have 5 rows of pixels because the FCM has 5 nodes and each row represents the time evolution of one node. The color of the node represents its value. A bright color corresponds to a high value. Yellow nodes have value 1. Purple nodes have value 0. The top figure starts with the initial state $\begin{pmatrix} 0 & 0 & 0 & 1 & 0 \end{pmatrix}$ and the bottom figure starts with the initial state $\begin{pmatrix} 0 & 1 & 0 & 1 & 0 \end{pmatrix}$. They both fall into the same limit cycle: $\begin{pmatrix} 0 & 0 & 0 & 1 & 0 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & 0 & 0 & 0 & 1 \end{pmatrix} \rightarrow \begin{pmatrix} 0 & 1 & 0 & 0 & 0 \end{pmatrix} \rightarrow \begin{pmatrix} 0 & 0 & 1 & 0 & 0 \end{pmatrix} \rightarrow \begin{pmatrix} 0 & 0 & 0 & 1 & 0 \end{pmatrix}$. The middle figure starts with initial state $\begin{pmatrix} 0 & 0 & 0 & 1 & 1 \end{pmatrix}$ and falls into a different limit cycle: $\begin{pmatrix} 1 & 0 & 0 & 1 & 0 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & 1 & 0 & 0 & 1 \end{pmatrix} \rightarrow \begin{pmatrix} 0 & 1 & 1 & 0 & 0 \end{pmatrix} \rightarrow \begin{pmatrix} 0 & 0 & 1 & 1 & 0 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & 0 & 0 & 1 & 0 \end{pmatrix}$.

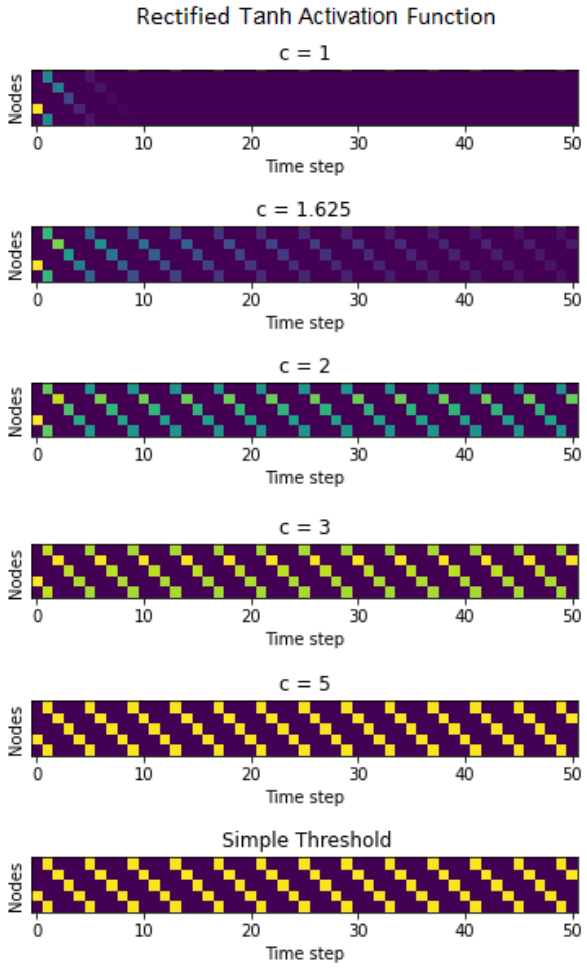

Fig. 5. Time evolution of the Dolphin FCM that uses rectified hyperbolic tangent (ReTanh) function Φ concept nodes. The time step is along the x-axis. The nodes are along the y-axis. The color of the node represents its value. Bright color represents high value. Yellow nodes have value 1. Purple nodes have value 0. The initial state is $\begin{pmatrix} 0 & 0 & 0 & 1 & 0 \end{pmatrix}$ and converges to a limit cycle in the ReTanh FCM. These limit cycles are similar to the limit cycles of the threshold FCM in Figure 4. The last image shows the limit cycle of the threshold FCM. The limit cycle decays for low values of c.

matrix E into a block matrix:

$$E = \begin{pmatrix} E_P & E_{PO} \\ E_{OP} & E_O \end{pmatrix} \tag{5}$$

where E_O contains all the edges between observable nodes, E_P contains all the edges between phantom nodes, E_{PO} contains all the edges from a phantom node to an observable one, and E_{OP} contains all the edges from an observable node to a phantom one.

Consider the Dolphin FCM. Suppose that an expert does not know how herd clustering affects dolphin behavior. So the herd-clustering node C_1 does not appear in the expert's FCM even though the node does affect the system dynamics. Then C_1 acts a phantom node as Figure 6 shows. This would give the terms E_O, E_P, E_{PO}, and E_{OP} as

$$E_O = \begin{pmatrix} 0 & 1 & 0 & -1 \\ -1 & 0 & 1 & -1 \\ 0 & -1 & 0 & 1 \\ 1 & 0 & -1 & 0 \end{pmatrix} \tag{6}$$

$$E_P = \begin{pmatrix} 0 \end{pmatrix} \tag{7}$$

$$E_{PO} = \begin{pmatrix} 1 & 0 & -1 & 0 \end{pmatrix} \tag{8}$$

$$E_{OP}^T = \begin{pmatrix} 0 & 0 & 1 & -1 \end{pmatrix} . \tag{9}$$

Consider next where the experts do not know or model

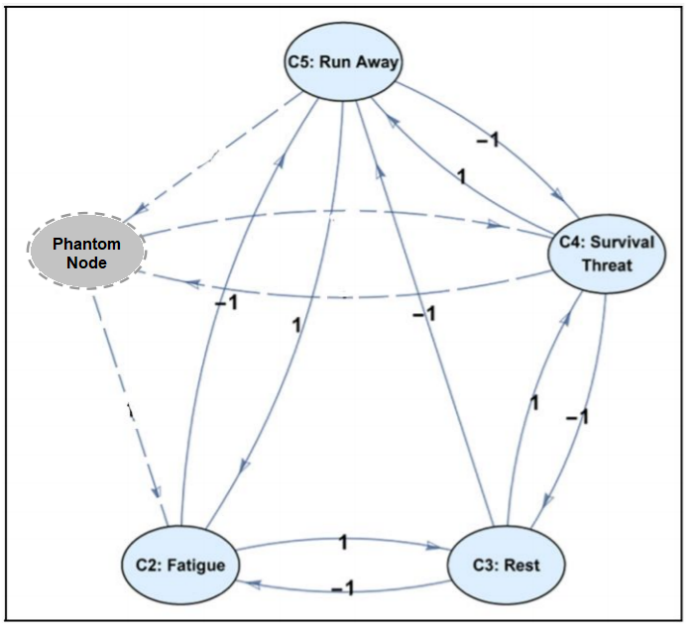

Fig. 6. The Dolphin FCM with a phantom node C_1 that corresponds to dolphin herd clustering. The domain expert or causal learning system does not observe C_1 or know the edges that correspond to C_1.

the predators that also live in the sea. Their 4-node FCM will not include the "SURVIVAL THREAT" node. This FCM predicts that the resting dolphins will keep resting because there are no predators. The actual behavior of the dolphins will cycle between resting from fatigue and running away in herd clusters. Explaining this cyclical behavior requires the phantom node "SURVIVAL THREAT" as in Figure 7.

The learning system should detect correct or accurate phantom nodes and not just any phantom nodes. So it needs to accurately estimate the causal edge degeess E_P, E_{OP}, and E_{PO}. Consider the Dolphin FCM with phantom node "SURVIVAL THREAT". Suppose an expert wrongly estimates the respective causal edges $E_{OP} = \begin{pmatrix} 1 & 0 & -1 & 1 \end{pmatrix}$ and $E_{PO} = \begin{pmatrix} -1 & 0 & 1 & -1 \end{pmatrix}$ as $\begin{pmatrix} 0 & 1 & 0 & 1 \end{pmatrix}$ and $\begin{pmatrix} 0 & 1 & 0 & 1 \end{pmatrix}$. Then the resulting phantom node predicts that the dolphins will rest in the presence of survival threats such as lethal sharks. But actual dolphin behavior would differ because the dolphins will flee and later rest in perhaps fall into a flee-rest cycle. Figure 8 shows this aberrant effect of a bad phantom node.

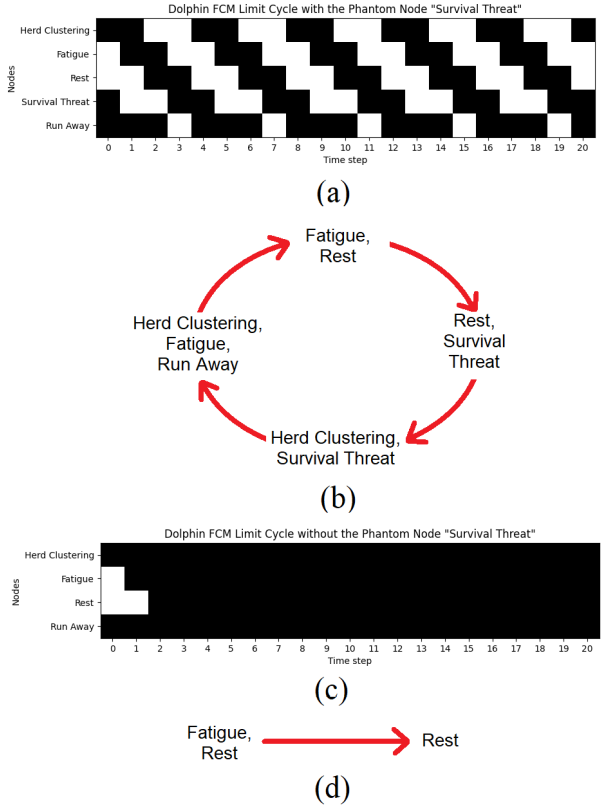

Fig. 7. Dolphin FCM limit cycles with and without phantom nodes. (a) Limit cycle of the Dolphin FCM with "SURVIVAL THREAT" as a phantom node. The FCM starts at the initial state $\begin{pmatrix}0 & 1 & 1 & 0 & 0\end{pmatrix}$ and then goes through the cycle $\begin{pmatrix}0 & 1 & 1 & 0 & 0\end{pmatrix} \rightarrow \begin{pmatrix}0 & 0 & 1 & 1 & 0\end{pmatrix} \rightarrow \begin{pmatrix}1 & 0 & 0 & 1 & 0\end{pmatrix} \rightarrow \begin{pmatrix}1 & 1 & 0 & 1 & 0\end{pmatrix} \rightarrow \begin{pmatrix}0 & 1 & 1 & 0 & 0\end{pmatrix}$. (b) This limit cycle corresponds to the cycle of the dolphins as they rest from fatigue, face a survival threat while resting, cluster in a herd to evade the threat, and get tired again from running away in a herd. (c) The limit cycle of the Dolphin FCM without the "SURVIVAL THREAT" phantom node. The FCM starts at the same initial state $\begin{pmatrix}0 & 1 & 1 & 0\end{pmatrix}$. But it goes through the state $\begin{pmatrix}0 & 0 & 1 & 0\end{pmatrix}$ and then gets stuck at the fixed-point state $\begin{pmatrix}0 & 0 & 0 & 0\end{pmatrix}$. (d) This corresponds to the dolphins staying at rest because there is no threat.

IV. Learning Phantom Edge-Weights from Limit-Cycles

The fixed points, limit cycles and chaotic attractors characterize a dynamical system $\dot{x} = f(x)$. Differential Hebbian learning can approximate the edge values from the node values on either side of the edge without sampling from the limit cycles. But we do not have samples from the phantom node. We can instead learn from the observed equilibrium attractors themselves. We can learn the phantom edges that approximate the observed fixed points and limit cycles if we augment the observed FCM.

Consider a FCM with n observable nodes. Let $C(t) \in [0,1]^n$ denote the state vector of the observable nodes at time t. Let $\hat{C}(t) \in [0,1]^n$ denote the vector of the observable node values at time t that the experts' FCM gives when augmented with the phantom nodes. The k-step squared error in the limit-cycle approximation is

$$L = \sum_{t=1}^{k} ||C(t) - \hat{C}(t)||_2^2 \tag{10}$$

if $||\ ||_2$ denotes the l^2 norm of the n-dimensional vector. Let Θ_P denote the edges of the phantom nodes present in E_P, E_{PO}, and E_{OP}. Then the optimal value of Θ_P is

$$\Theta_P^* = \underset{\Theta_P}{\operatorname{argmin}} \sum_{t=1}^{k} ||C(t) - \hat{C}(t)||_2^2 \tag{11}$$

Gradient algorithms can give locally optimal values of Θ_P^*. Figure 9 shows this causal learning method based on minimizing the error between target and augmented limit cycles.

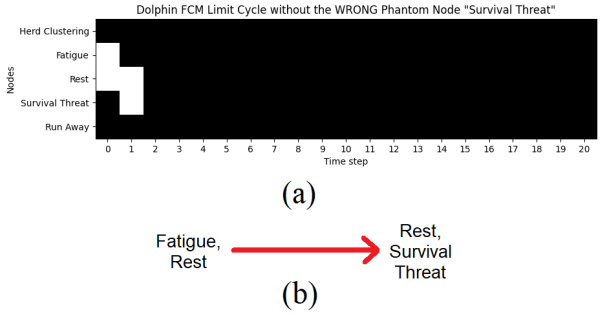

Fig. 8. (a) A limit cycle of the Dolphin FCM with a "bad" phantom node for "SURVIVAL THREAT". The FCM starts at the same initial state $\begin{pmatrix}0 & 1 & 1 & 0 & 0\end{pmatrix}$. But it goes through the state $\begin{pmatrix}0 & 0 & 1 & 1 & 0\end{pmatrix}$ and then gets stuck at the fixed point state $\begin{pmatrix}0 & 0 & 0 & 0 & 0\end{pmatrix}$. (b) This corresponds to the dolphins staying at rest even in the presence of a "SURVIVAL THREAT" such as a shark.

V. Experimental Results

We ran the Dolphin FCM with the 4 observable nodes C_2, C_3, C_4, and C_5 and one phantom node C_1 using 10,000 random initial states. The phantom-edge approximator trained on the node vector $C(t)$ for the first k = 2 steps of each FCM run and approximated Θ_P with Θ_P^*. This Θ_P^* augmented the observable edge matrix E_O to give

$$E^* = \begin{pmatrix} E_P^* & E_{PO}^* \\ E_{OP}^* & E_O \end{pmatrix} \tag{12}$$

$$= \begin{pmatrix} 0 & 0.6685 & 0.4392 & 0.0066 & 0.8296 \\ 0.6685 & 0 & 1 & 0 & -1 \\ 0.4392 & -1 & 0 & 1 & -1 \\ 0.0066 & 0 & -1 & 0 & 1 \\ 0.8296 & 1 & 0 & -1 & 0 \end{pmatrix} \tag{13}$$

This edge matrix E^* differs somewhat from the original Dolphin edge matrix E in (4). But their equilibrium dynamics are similar because E^* gives a FCM that had similar limit cycles to the original Dolphin FCM. Figure 10 shows an instance of this causal learning.

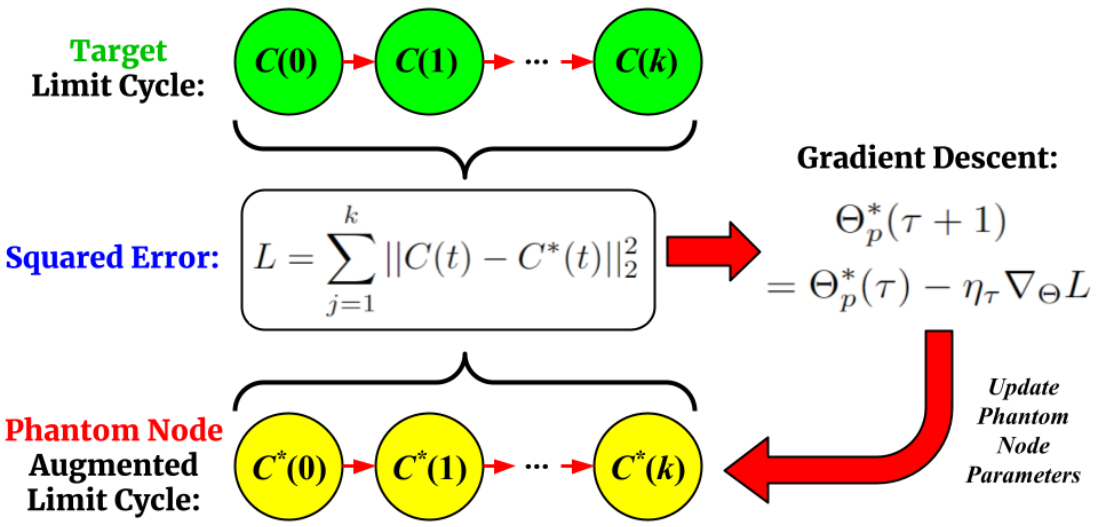

Fig. 9. Phantom node learning through gradient descent on sampled limit-cycle equilibria. The gradient-based updates minimize the squared error between the target limit cycle and the augmented FCM's limit cycle over k consecutive steps.

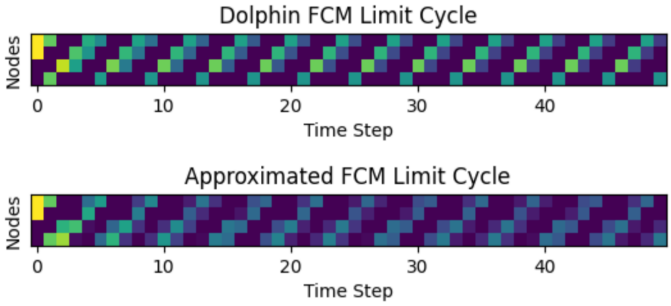

Fig. 10. The approximated limit cycles through phantom-node causal learning. The time step is along the x-axis. The 4 observable nodes lie along the y-axis. The color of the node represents its value. Bright color represents high value. Yellow nodes have value 1. Purple nodes have value 0. The initial state is $\begin{pmatrix} 1 & 1 & 0 & 0 \end{pmatrix}$ and converges to limit cycles in both FCMs. These approximated limit cycles are similar to the limit cycles of the Dolphin FCM.

VI. Conclusion

Detecting and absorbing phantom nodes is a difficult problem in causal modeling. New causal nodes change the equilibrium outcomes of feedback causal models such as FCMs. This paper shows one way to learn a phantom node's causal edge structure in the supervised case where the user knows or estimates the equilibrium limit cycles of the full FCM. Learning phantom nodes in this case may give some confidence in time-series training based on differential Hebbian or other unsupervised learning schemes. Many other supervised and unsupervised learning schemes can apply to learn one or more phantom nodes given some knowledge of the causal edges or of the limit cycles. These schemes can add extra computational cost when the number of phantom nodes increases.

References

[1] Vilayanur S Ramachandran and Diane Rogers-Ramachandran. Phantom limbs and neural plasticity. *Archives of neurology*, 57(3):317–320, 2000.

[2] Bart Kosko. Fuzzy cognitive maps. *International journal of man-machine studies*, 24(1):65–75, 1986.

[3] Bart Kosko. Hidden patterns in combined and adaptive knowledge networks. *International Journal of Approximate Reasoning*, 2(4):377–393, 1988.

[4] Hussam Alkaissi and Samy I McFarlane. Artificial hallucinations in ChatGPT: implications in scientific writing. *Cureus*, 15(2), 2023.

[5] Bart Kosko and Sanya Mitaim. Stochastic resonance in noisy threshold neurons. *Neural networks*, 16(5-6):755–761, 2003.

[6] Patrik O Hoyer, Shohei Shimizu, Antti J Kerminen, and Markus Palviainen. Estimation of causal effects using linear non-gaussian causal models with hidden variables. *International Journal of Approximate Reasoning*, 49(2):362–378, 2008.

[7] Sosuke Ito. Backward transfer entropy: Informational measure for detecting hidden markov models and its interpretations in thermodynamics, gambling and causality. *Scientific reports*, 6(1):36831, 2016.

[8] Wei Chen, Ruichu Cai, Zhifeng Hao, Chang Yuan, and Feng Xie. Mining hidden non-redundant causal relationships in online social networks. *Neural Computing and Applications*, 32:6913–6923, 2020.

[9] Bart Kosko. Differential Hebbian learning. In *AIP Conference proceedings*, volume 151, pages 277–282. American Institute of Physics, 1986.

[10] Osonde A Osoba and Bart Kosko. Fuzzy cognitive maps of public support for insurgency and terrorism. *The Journal of Defense Modeling and Simulation*, 14(1):17–32, 2017.

[11] Guy Ziv, Elizabeth Watson, Dylan Young, David C Howard, Shaun T Larcom, and Andrew J Tanentzap. The potential impact of brexit on the energy, water and food nexus in the uk: A fuzzy cognitive mapping approach. *Applied Energy*, 210:487–498, 2018.

[12] Michael Glykas. *Fuzzy cognitive maps: Advances in theory, methodologies, tools and applications*, volume 247. Springer, 2010.

[13] Elpiniki I Papageorgiou. *Fuzzy cognitive maps for applied sciences and engineering: from fundamentals to extensions and learning algorithms*, volume 54. Springer Science & Business Media, 2013.

[14] Wojciech Stach, Lukasz Kurgan, and Witold Pedrycz. A divide and conquer method for learning large fuzzy cognitive maps. *Fuzzy Sets and Systems*, 161(19):2515–2532, 2010.

[15] Rod Taber, Ronald R Yager, and Cathy M Helgason. Quantization effects on the equilibrium behavior of combined fuzzy cognitive maps. *International Journal of Intelligent Systems*, 22(2):181–202, 2007.

[16] Julie A Dickerson and Bart Kosko. Virtual worlds as fuzzy cognitive maps. *Presence: Teleoperators & Virtual Environments*, 3(2):173–189, 1994.

[17] Judea Pearl. *Causality*. Cambridge University Press, 2009.