

Converting Sequenced Fuzzy Cognitive Maps to Causal Virtual Worlds with Large Video Generators

Akash Kumar Panda¹ and Olaoluwa Adigun² and Bart Kosko³

Abstract—We show how users can create and manipulate causal virtual worlds with large-language-model (LLM) and large-video-model agents. The approach uses feedback fuzzy cognitive maps (FCMs) both to model the granular causal structure of the virtual world and to guide its causal evolution. The local causal rules are partial or fuzzy while the FCM’s feedback structure produces global equilibria that define causal scenarios. A sequence of *dynamical* meta-rules of the form “If $\mathcal{A}$ then $\mathcal{B}$ ” define the causal scenes of the virtual-world video. The if-part causal pattern $\mathcal{A}$ perturbs the FCM’s virtual world at the user’s or agent’s discretion. The FCM’s transient feedback dynamics define the meta-rule’s causal arrow of implication. The then-part $\mathcal{B}$ is the resulting equilibrium attractor such as a FCM limit cycle or fixed point. Our algorithm extracts these meta-rules from the FCM and guides the LLM agent to write a script based on the FCM meta-rule sequence. The large video generator converts the meta-rule into a video scene in accord with the flow of the dynamics. We applied the agent-based technique to a simple FCM that describes an undersea world of dolphins and sharks. Google’s Gemini 3.1 generated the script and Google’s Veo 3.1 generated the dolphin-shark video. The approach is general and can scale by mixing larger FCMs and AI agents to produce more immersive virtual worlds.

I. AGENTIC AI GENERATION OF VIRTUAL WORLDS FROM SEQUENCED FUZZY COGNITIVE MAPS

We show how to generate realistic causal virtual-world videos by combining feedback Fuzzy Cognitive Maps (FCMs) with large-video-generator agents.

This new virtual-world technique is general and scales for arbitrary mixed causal-graph FCMs combined with controlled AI agents. As OpenAI has stated: “scaling video generation models is a promising path towards building general purpose simulators of the physical world.” [1] Sequenced FCM dynamics offers a practical way to approximate such models.

Figure 1 shows the system flow of stimulated FCM causal patterns that produce an AI-generated 1-minute video from a simple 5-node FCM. The FCM defines a coarse-grained undersea world of dolphins and sharks. We use this simple dolphin FCM because over the decades it has become a type of small-scale test and teaching model of an FCM [2]. The FCM itself is a nonlinear feedback dynamical system whose causal graph is signed and weighted or partial to reflect

partial or fuzzy causality [3], [4], [4]–[10]. The local causal rule structure of the FCM helps make the nonlinear system interpretable or explainable [11], [12].

The FCM-generated video shows how a pod of playful dolphins responds to the threat of a large tiger shark. The video uses a sequence of 8 scenes as we explain in detail in the next two sections. Figure 2 shows further snapshots from the video.

Figure 3 shows the 5-node dolphin-shark FCM and its corresponding 5-by-5 causal edge matrix E . This paper uses just this small FCM to illustrate the process of generating virtual worlds although we stress again that the process is completely general for any size or mixture of FCMs.

Figure 4 shows the sequence of 8 scenes as 8 input-output pairs of stimulated FCM states that map to equilibrium limit cycles or fixed-point attractors. This gives a type of dynamical system “storyboard” of the FCM-based video or virtual world. Users can pick the stimulus patterns directly or ask the AI agent to pick the nearest-matching patterns for them. Users can also allow an agent to pick the stimulus patterns itself.

Figures 5-8 show frames from the first 4 scenes of the final video. Figure 9 shows all sequenced limit cycles used.

The key insight is that a FCM dynamical system defines a set of causal *meta-rules* of the form $\mathcal{A}_j \rightarrow \mathcal{B}_j$. We “sequence” or impose a time order on these FCM meta-rules. We then use the sequenced meta-rules to drive the time-ordered scenes of the generated video or virtual world.

The nonlinear FCM dynamical system $\mathcal{F}$ itself is a time-varying mapping $\mathcal{F} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ from a causal pattern space such as $\mathbb{R}^n$ back into itself. A pulsed or sustained (“clamped”) input vector $x \in \mathbb{R}^n$ produces a sequence of transient causal states that ends in an equilibrium attractor that can range from a fixed point to a limit cycle to aperiodic chaos. The simple binary-state dolphin FCM in Figure 3 always ends in a limit cycle or a fixed-point attractor that lives inside the binary 5-cube $\{0, 1\}^5$. The fuzzy causal edge values e_{ij} of the directed causal edge from concept node C_i to node C_j takes values in the bipolar interval $[-1, 1]$ in general where positive values $e_{ij} > 0$ denote causal increase and negative values $e_{ij} < 0$ denote causal decrease and $e_{ij} = 0$ denotes no causal effect.

The global meta-rule $\mathcal{A}_j \rightarrow \mathcal{B}_j$ asks and answers its own “what-if” question: What dynamical equilibrium $\mathcal{B}_j$ results if we perturb the FCM causal system with new input x or otherwise knock it into a new state of *disequilibrium* $\mathcal{A}_j$?

So the implication arrow “ $\rightarrow$ ” in the meta-rule $\mathcal{A}_j \rightarrow \mathcal{B}_j$ is not a classical logical or material-implication operator. The

¹Akash Kumar Panda is with the Department of Electrical and Computer Engineering, University of Southern California, Los Angeles, CA, USA akashpan@usc.edu

²Olaoluwa Adigun is with the School of Computing and Information Sciences, Florida International University Miami, FL 33199, USA olaadigu@fiu.edu

³Bart Kosko is with the Department of Electrical and Computer Engineering, University of Southern California, Los Angeles, CA, USA kosko@usc.edu

arrow is instead a *causal implicant* that depends on system transients and hence depends on time in general. These are unconstrained rules. The user can add *constraint* rules that guide or extend the equilibria and so modify the FCM video.

The resulting set of such dynamical meta-rules $\mathcal{A}_j \rightarrow \mathcal{B}_j$ endows the FCM system with a global level of explainability or XAI. This still holds in the general case when we take convex combinations $w_1\mathcal{F}_1 + \dots + w_K\mathcal{F}_K$ of K -many FCMs $\mathcal{F}_1, \dots, \mathcal{F}_K$ for nonnegative mixing weights w_j that sum to unity and that can depend on x . Mixing FCM causal edge matrices $E_1, \dots, E_K$ always produces a new causal edge matrix and thus always produces an FCM $\mathcal{F}$. Appropriate zero-padding of rows and columns ensures that the mixed causal edge matrices are conformable for matrix addition.

This mixing-closure property greatly extends the application reach and knowledge-representation power of FCMs. Users can combine and uncombine and adapt the underlying mixed FCM at will. Users can also pick the mixed meta rules and especially the disequilibrium stimuli $\mathcal{A}_j$ at will to create and change the virtual world. Agentic AI systems can further assist in this immersive world-creating process.

The next section reviews FCMs and show how to sequence and fire their meta-rules. It shows how the dolphin FCM's 8 causal meta-rules $\mathcal{A}_1 \rightarrow \mathcal{B}_1, \dots, \mathcal{A}_8 \rightarrow \mathcal{B}_8$ become the 8 scene prompts to the AI video generator.

The last section explains how to convert these FCM scenes to prompts that control Google's Veo 3.1 large-video generator by way of Algorithm 1 below. We here focus on user-driven choices of meta-rules $\mathcal{A}_j \rightarrow \mathcal{B}_j$ and resulting scenes but recent work shows that we can easily use large language models (LLMs) to generate FCM knowledge graphs from text or vice versa [13] in more complex agentic AI systems.

II. SEQUENCED FUZZY COGNITIVE MAPS

A. Fuzzy Cognitive Maps (FCM)

FCMs model causal dynamical systems as directed weighted graphs [3], [4], [4]–[10]. The concept nodes describe the causal variables in the dynamical system and the directed edges describe the causal relationships between those nodes. FCMs allow feedback and therefore converge to non-trivial equilibria like limit cycles. FCMs model dynamical systems by approximating their underlying maps from inputs to equilibria.

B. Causal Edge Matrix

The directed edges of the FCM describe the causal relationships between concept nodes. An edge e_{ij} from the i^{th} concept node C_i to the j^{th} concept node C_j means “ C_i causes C_j ”. The fuzzy edge weights describe partial causality and we use the same symbol for this value—it can also vary with time. The causal edge weight $e_{ij} \in [-1, 1]$ on the ij th edge states the degree to which C_i causes C_j :

$$e_{ij} = \text{Degree}(C_i \rightarrow C_j). \quad (1)$$

A positive e_{ij} means that C_j increases when C_i increases and a negative e_{ij} means that C_j decreases if C_i increases. The weight e_{ij} is zero when there is no causal edge between

C_i and C_j . The magnitude of e_{ij} is high if there is a strong causal relationship between C_i and C_j . A low magnitude of e_{ij} describes a weak causal relationship between C_i and C_j .

An $n \times n$ matrix E describes all the directed weighted edges of a n -node FCM. The causal edge weight e_{ij} corresponds to the matrix element on the i^{th} row and the j^{th} column. The matrix element is zero if there is no edge between the corresponding node-pair.

Figure 3 shows the dolphin FCM with 5 nodes: “Herd Clustering”, “Fatigue”, “Rest”, “Survival Threat”, and “Run Away”. The figure also shows its corresponding causal edge matrix. The matrix element e_{45} on the 4th row and the 5th column is 1 because there is a positive edge with weight 1 from the 4th concept node “Survival Threat” to the 5th concept node “Run Away”. This says that the dolphins “run away” when a “survival threat” like a shark is present.

C. FCM Evolution

A n -dimensional row vector $C(t) \in [0, 1]^n$ describes the state of the FCM's concept nodes at time t . The i^{th} node is “active” or “on” at time t if the i^{th} component $C_i(t)$ of the state vector $C(t)$ is equal to or close to one. The i^{th} node is “inactive” or “off” at time t if $C_i(t)$ is equal to or close to zero. A node is partially active otherwise. The causal variables corresponding to the active nodes are present in the system and those corresponding to the inactive nodes are absent. The causal factors are partially present in the system if their corresponding nodes are partially active.

Consider the dolphin FCM in figure 3. If $C(t) = (0 \ 0 \ 0 \ 1 \ 0)$ then the 4th node “Survival Threat” is “on” and all other nodes are “off” at time t . This means there is a “threat” like a shark lurking around.

FCMs evolve in discrete time through vector-matrix multiplication and nonlinear compression—at least in the simple FCM we use for this video creation. The state $C_j(t+1)$ of the j^{th} concept node C_j at discrete time step $t+1$ is:

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

if Φ is a nonlinear function bounded between zero and one.

The sum $\sum_{i=1}^n C_i(t) e_{ij}$ is the matrix product between the state row-vector $C(t)$ and the edge matrix E . The nonlinear function Φ then compresses this product between zero and one. We stress that Φ is an arbitrary but bounded nonlinear function in general even though here we use the simplest case of a threshold nonlinearity.

This process repeats itself to give the discrete-time evolution of the FCM. The FCM starts with the initial state $C(0)$ at time $t=0$ and then goes through the states $C(1)$, $C(2)$, $C(3)$, and so on in order. The active nodes in this state-vector sequence qualitatively describe the trajectory of the dynamical system that the FCM models.

Consider the dolphin FCM with an initial state $C(0) = (0 \ 0 \ 0 \ 1 \ 0)$. Equation 2 and figure 3(b) give $C(1) = (1 \ 0 \ 0 \ 0 \ 1)$. “Survival Threat” was active at $t=0$ and “Herd Clustering” and “Run Away” were active at $t=1$. This shows the dolphins' behavior in the presence of a




Fig. 1: Frames from a video of the virtual world generated from a 5-node Fuzzy Cognitive Map (FCM) that models dolphin behavior in presence of a “Survival Threat”. The video shows a pod of dolphins running away from a shark over 8 scenes. The input stimulus for each scene is above the frame corresponding to the scene. The figures on the left and right of the video frames show the respective state of the FCM at the beginning and the end of the video. The figure highlights clamped node “Herd Clustering” in green.

shark. They group themselves together in herd clusters and run away from the shark.

D. FCM Equilibria

The equilibria characterize a dynamical system. The equilibrium behavior of the FCM depends on the limiting behavior of the state-vector sequence. The FCM converges to a “fixed point” if the state-vector sequence converges to a constant vector. The FCM converges to a K -step “limit cycle” for an integer $K > 1$ if $C(t+K) = C(t)$ somewhere in the state-vector sequence. Then the FCM converges to an equilibrium where K state vectors repeat themselves over and over in the same order. The FCM may also converge to a chaotic attractor where there are no repeating patterns in the state-vector sequence.

Consider the dolphin FCM with same initial state $C(0) = (0 \ 0 \ 0 \ 1 \ 0)$. Equation 2 and figure 3(b) give $C(1) = (1 \ 0 \ 0 \ 0 \ 1)$, $C(2) = (0 \ 1 \ 0 \ 0 \ 0)$, $C(3) = (0 \ 0 \ 1 \ 0 \ 0)$, and $C(4) = (0 \ 0 \ 0 \ 1 \ 0) = C(0)$. The state vectors $C(0)$, $C(1)$, $C(2)$, and $C(3)$ repeat themselves over and over in the same order. The dolphin FCM converges to a 4-state limit cycle. The active nodes in these states respectively are “Survival Threat”, “Herd Clustering” and “Run Away”, “Fatigue”, and “Rest”. So dolphins run away from threats like sharks in herd clusters. They then get tired and need to rest. The shark then catches up to the dolphins while they rest. This starts the cycle all over again.

The set of all initial conditions $C(0)$ that lead to a given equilibrium describes the “basin of attraction” for that equilibrium. The FCM describes a map from these basins to their corresponding equilibrium attractors. The basins

of the FCM’s equilibria partition the FCM’s input space. The FCM models a dynamical system by approximating its corresponding basin-to-equilibrium map.

E. Causal State Activation: Clamping vs. Pulsing

“Clamping” the k^{th} node C_k to a value $c_k \in [0, 1]$ forces $C_k(t)$ to equal c_k irrespective of the values of other nodes at time $t - 1$. The k^{th} node is clamped “on” if $c_k \approx 1$ and it is clamped “off” if $c_k \approx 0$. Clamping acts as a forcing function and pushes the FCM into new equilibria. FCMs can clamp one or more nodes at a time. Clamping is often a way to implement policies and answer What-if questions regarding their effects.

Clamping on the “Survival Threat” node changes the dolphin FCM’s limit cycle from $(0 \ 0 \ 0 \ 1 \ 0) \rightarrow (1 \ 0 \ 0 \ 0 \ 1) \rightarrow (0 \ 1 \ 0 \ 0 \ 0) \rightarrow (0 \ 0 \ 1 \ 0 \ 0) \rightarrow (0 \ 0 \ 0 \ 1 \ 0)$ to a new 3-state limit cycle $(0 \ 1 \ 0 \ 1 \ 0) \rightarrow (1 \ 0 \ 0 \ 1 \ 0) \rightarrow (1 \ 1 \ 0 \ 1 \ 0) \rightarrow (0 \ 1 \ 0 \ 1 \ 0)$. The active nodes in this new limit cycles respectively are “Fatigue” and “Survival Threat”; “Herd Clustering” and “Survival Threat”; and “Herd Clustering”, “Fatigue”, and “Survival Threat”. The shark is *always* present in this case. The dolphins cluster and get tired over and over. The nodes “Rest” and “Run Away” always off because the dolphins do not get a chance to rest and they never escape the shark.

FCMs can also ‘pulse’ a node instead of clamping it. ‘Pulsing’ the k^{th} node C_k to a value c_k at time t_0 forces $C_k(t_0)$ to equal c_k irrespective of the values of other nodes at time $t_0 - 1$. This only affects the FCM at time t_0 and the FCM reverts back to its unperturbed dynamics after t_0 . The



(a) Scene 1



(b) Scene 2



(c) Scene 3



(d) Scene 4



(e) Scene 5



(f) Scene 6




(g) Scene 7



(h) Scene 8

Fig. 2: Limit cycle 1: Interaction of a dolphin pod and a tiger shark in a water. (a). Scene 1: The dolphin pod rests and plays. (b). Scene 2: The dolphin pod runs away from a threat, gets tired, and rests. (c). Scene 3: The dolphins swim in a herd cluster to avoid a large tiger shark. (d). Scene 4: The tiger shark breaks the dolphins' herd-cluster apart as the dolphins tire. (e). Scene 5: The tired dolphins stop running away and regroup in a tighter herd-cluster. (f). Scene 6: The dolphins push through their fatigue and swim away from the large tiger shark. (g). Scene 7: The dolphin pod gets away from the large tiger shark. (h). Scene 8: The dolphin goes back to playing and resting until their next shark encounter.



(a)

$$E = \begin{matrix} & \begin{matrix} C_1 & C_2 & C_3 & C_4 & C_5 \end{matrix} \\ \begin{matrix} C_1 \\ C_2 \\ C_3 \\ C_4 \\ C_5 \end{matrix} & \begin{bmatrix} 0 & 1 & 0 & -1 & 0 \\ 0 & 0 & 1 & 0 & -1 \\ 0 & -1 & 0 & 1 & -1 \\ 1 & 0 & -1 & 0 & 1 \\ -1 & 1 & 0 & -1 & 0 \end{bmatrix} \end{matrix}$$

(b)

Fig. 3: The 5-node Dolphin FCM with trivalent causal edges. (a) The directed weighted graph for the dolphin FCM. The directed edge e_{ij} is trivalent because $e_{ij} \in \{-1, 0, 1\}$ but in general can be properly fuzzy and take on any directed causal degree in the bipolar interval $[-1, 1]$. The positive edges are in blue and the negative edges are in red. (b) The causal edge matrix E that corresponds to the dolphin FCM. The 5 concept nodes C_1 - C_5 index both the rows and the columns. The matrix elements give the edge weights e_{ij} . The positive edge values $e_{ij} > 0$ are in blue and the negative edge values $e_{ij} < 0$ are in red.

node C_k pulses ‘on’ at time t_0 if $C_k(t_0) \approx 1$ and it pulses ‘off’ at time t_0 if $C_k(t_0) \approx 0$. FCMs can also pulse more than one node at a time.

The unperturbed dolphin FCM may converge to its fixed point $(0 \ 0 \ 0 \ 0 \ 0)$. “Survival Threat” ‘on’-pulse can drive the FCM into a 4-stage limit cycle instead.

F. FCM Sequencing

FCM sequencing stimulates the FCM with a sequence of clamping patterns $\mathcal{A}_k$ to push it through a corresponding sequence of equilibria $\mathcal{B}_k$. This map from stimulus patterns $\mathcal{A}_k$ to their corresponding FCM equilibria $\mathcal{B}_k$ define a set of causal ‘if-then’ meta-rules $\mathcal{A}_k \rightarrow \mathcal{B}_k$. The causal if-part $\mathcal{A}_k$ stimulates the FCM at the user’s discretion through clamps and pulses. The arrow $\rightarrow$ describes the transient dynamics that connect $\mathcal{A}_k$ to its corresponding equilibrium $\mathcal{B}_k$. The stimulus pattern destroys the FCM’s prior equilibrium and before pushing it into a new one.

Figure 4 shows this process for a sequence of 8 if-parts $\mathcal{A}_1$ - $\mathcal{A}_8$ to the dolphin FCM. The figure also shows the corresponding then-parts $\mathcal{B}_1$ - $\mathcal{B}_8$. The transition from the if-part stimulus to the then-part equilibrium defines each of the 8 scenes in the virtual world video.

Scene 1 clamps on “Herd Clustering” and shows the FCM converging to a 4-stage limit cycle. The dolphins play

Algorithm 1: Causal Virtual World Generation

Input: $\mathcal{G}=(\mathcal{V}, \mathbf{E})$: FCM with $\mathbf{E} \in \mathbb{R}^{n \times n}$;
 $\mathbf{c}^{(0)} \in \{0, 1\}^n$: initial state and clamping conditions; T : time steps; $\varphi(\cdot)$: threshold function; $\mathcal{P}_{\text{script}}$: LLM scripting prompt; $\mathcal{P}_{\text{video}}$: LVM generation prompt; $\text{LLM}(\cdot)$, $\text{LVM}(\cdot)$: language and video models

Output: $\mathcal{D}$: causal virtual world documentary

// Step 1: FCM Simulation

```

1 Initialize  $\mathbf{C} \leftarrow \{ \}$ ,  $\mathbf{c} \leftarrow \mathbf{c}^{(0)}$ ;
2 for  $t = 0$  to  $T - 1$  do
3   | Store  $\mathbf{C}[\cdot, t] \leftarrow \mathbf{c}$ , then update  $\mathbf{c} \leftarrow \varphi(\mathbf{E} \cdot \mathbf{c})$ ;
4 end
5 Detect equilibrium: set  $eq \leftarrow \text{FixedPoint}$  or  $\text{LimitCycle}$ ;

// Step 2: LLM Script Generation
6 Analyze  $\mathbf{C}$  for START, STOP, PERSIST transitions across all  $t$ ;
7  $SC \leftarrow \text{LLM}(\mathcal{P}_{\text{script}}, \mathcal{V}, \mathbf{C}, eq)$ ; // yields title +  $T$  scene descriptors  $\{p_t, lt_t\}$ 

// Step 3: LVM Clip Generation and Assembly
8  $\mathcal{D} \leftarrow \text{LVM}(\mathcal{P}_{\text{video}}, \{p_t\}, \{lt_t\}, SC.title)$ ;
// generate and concatenate clips
9 return  $\mathcal{D}$ ;
```

together as a group, get tired, rest, and then go back to playing. Scene 2 disrupts this equilibrium $\mathcal{B}_1$ by unclamping “Herd Clustering” and ‘on’-pulsing “Survival Threat”. This moves the system from $\mathcal{B}_1$ to another 4-stage limit cycle $\mathcal{B}_2$ where the dolphins come across a shark, run away, get tired, rest, and encounter the shark again. Scene 3 then interrupts this equilibrium through stimulus $\mathcal{A}_3$. This process then repeats 5 more times with stimuli $\mathcal{A}_4$ - $\mathcal{A}_8$. The FCM in turn converges to 2 more limit cycles $\mathcal{B}_4$ and $\mathcal{B}_8$ and 3 more fixed points $\mathcal{B}_5$ - $\mathcal{B}_7$. The 8 scenes together tell the story of a pod of dolphins that gets chased by a persistent shark and has to push through fatigue to escape.

Figures 5-8 show the equilibria $\mathcal{B}_1$ - $\mathcal{B}_4$ in detail. They show one frame from each state of the limit cycles. They also show the active nodes in those states. The limit cycle for $\mathcal{B}_8$ is the same as $\mathcal{B}_2$. The FCM equilibria $\mathcal{B}_5$ - $\mathcal{B}_7$ are fixed points and thus have only one state. Figure 2(e)-(g) show the video frames that best represent those fixed-point states.

III. AGENTIC CONTROL OF LARGE VIDEO GENERATOR SCENE PROMPTS

We present the agentic control for generating the causal virtual world for an FCM. This explains how to generate a video from a sequence of equilibria for an FCM with a large video generator. The sequence of equilibria comes from the concatenation of equilibria that result from the combination of input pulse signals and node clamping operation. We tested our method with Google’s Veo 3.1 video generator.




Fig. 7: Scene 3 from the virtual world video. $\mathcal{A}_3$ stimulates the dolphin FCM by clamping on the “Survival Threat” node. The dolphin FCM then converges to a 3-stage limit cycle $\mathcal{B}_3$. The figure shows one frame from each stage of the limit cycle and also mentions their corresponding active nodes next to the frame.




Fig. 8: Scene 4 from the virtual world video. $\mathcal{A}_4$ stimulates the dolphin FCM by clamping on the “Survival Threat” node as well as clamping off the “Herd Clustering” node. The dolphin FCM then converges to a 4-stage limit cycle $\mathcal{B}_4$. The figure shows one frame from each stage of the limit cycle and also mentions their corresponding active nodes next to the frame.

screenwriting and movie production. But each method varies in autonomy level during the screenwriting step.

A. Semi-autonomous agentic control

An LLM agent executes the screenwriting phase of this agentic control method. This operator converts a sequence of equilibria into a script. Script scenes align with state transitions within the equilibria sequence. Each scene details granular information derived from active state variables or nodes. The agent reads the vector sequence directly or employs an intermediate visualization. The agent inputs the script into the video generator agent. This agent transforms the script into a video or storyboard.

Let us consider the 5-node dolphin FCM in Figure 3.

Figure 4 shows the links between a sequence of the dolphin FCM equilibria with pulsing and clamping. The sequence comprises of 8 equilibrium cases from the FCM. Figure 9 visualizes the 8 equilibria and their transitions. Each scene corresponds to an equilibrium state.

An LLM agent annotated the sequenced equilibria into a movie script. The movie script is divided into three parts: task description, summary, and the scene description for each equilibrium state. The task description captures the high-level description of the movie. The summary provides more specific information. The scene descriptions translates the state variables into visual elements, dialogue and character interactions that advance the central narrative as laid out in the summary.

Each scene in Figure 9 captures the causal transition from the causal if-part $\mathcal{A}_k$ to its corresponding equilibrium $\mathcal{B}_k$. Consider the first pair $(\mathcal{A}_1, \mathcal{B}_1)$. The first column $\mathbf{C}(0) = [C_1(0), C_2(0), C_3(0), C_4(0), C_5(0)]$. The *if-part* corresponds to clamping on the “Herd Clustering” node while other nodes start in the OFF state. The description of Scene 1 follow from the causal interaction from $\mathbf{C}(0)$ to $\mathbf{C}(3)$.

Here is a snippet of our prompt to the LLM for script generation.

A snippet from the script-generation prompt $\mathcal{P}_{\text{script}}$

Task Guidelines:

1. Analyze the Data:

- (a) **Step $t = 0$ (The Opening):** Establish the scene based on which nodes are active (1) in the first column.
- (b) **Transitions ($t = 1$ to $t = T$):** Compare column t with column $t - 1$.
 - (i) Identify what **Started** (0 to 1).
 - (ii) Identify what **Stopped** (1 to 0).
 - (iii) Identify what **Persisted** (1 to 1).
- (c) **Interpret the causal logic:** If Node A is active at $t = 1$ and Node B activates at $t = 2$, imply that A might be causing B.

2. Documentary Structure:

- (a) **Title:** Generate a thematic title based on the Node List concepts.
- (b) **Lower Thirds:** Provide short text overlays summarizing the state logic (e.g., “Rainfall triggers River Surge”).
- (c) **Narrative (Voiceover):** A 1-2 sentence script connecting the current state to the previous one.
- (d) **Conclusion:** Explicitly address the “Equilibrium Behavior” in the final scene.

Google’s Veo 3.1 generative agent converted the LLM-generated script into a movie. The maximum allowed duration for a Veo 3.1 generated video is 8 seconds. We used the multiple-shot approach to generate the video as a sequence of shots. Here is a snippet of the instruction we use for generating the video:




Fig. 9: Sequence of all FCM equilibria in the video: Clamped-on states are in red and clamped-off states are in blue.

An snippet of the Veo 3.1 instruction $\mathcal{P}_{\text{video}}$

Style: Cinematic 4K, underwater photography, realistic Pacific Bottlenose dolphins, deep ocean blue color grade, naturalistic lighting, fluid camera movements, high-fidelity water physics.

Prompt: Generate A pod of dolphins is seen from a side-angle tracking shot, swimming at maximum speed through the deep blue ocean. Intense cavitation bubbles trail from their fins. The camera moves rapidly alongside them. The scene ends with the dolphins beginning to slow down.

Transition Note: End shot with decelerating tail beats.

We sequentially prompted the video generator with the scene descriptions from the LLM annotated movie script. The agent rendered the Scene 1 description as an 8-second clip. The model extended the output to 16 seconds using the Scene 2 description. This recursive extension continued through all 8 scene descriptions. The procedure yielded a 57-second video encapsulating all scenes.

Figure 2 shows frames from the generated video. Figure 2a shows a frame from Scene 1 and corresponds to its equilibrium state in Figure 9. This scene runs from $t = 0$ to $t = 3$ with “Herd Clustering” node clamped ON and the “Survival Threat” node inactive: dolphins play and rest without a survival threat. Here is the annotation for this scene: *A pod of dolphin is resting and playing*. Figure 2a depicts this. Figure 2b - 2h similarly capture the corresponding causal virtual world of Scene 2-8 in Figure 9.

Again we can use LLM agents to extract FCMs from text or vice versa [13], [14] and exploit the Bayesian *de-chunking* structure [15]. This allows AI agents to partially or fully automate the selection of scenes or meta-rules $\mathcal{A}_j \rightarrow \mathcal{B}_j$.

IV. CONCLUSIONS

Sequenced FCMs can combine with agentic control of large video generators to produce causal virtual worlds or videos. Users can define or modify the FCMs or oversee LLMs that generate them from text or from transcribed speech. Users can also modify the virtual world midstream if they impose new causal patterns or if they adapt the FCM’s local causal structure. These techniques scale and favor hardware implementation both because FCM causal propagation involves only simple vector-matrix multiplication of often sparse causal edge matrices and because mixing

FCMs always produces a new FCM. Large-scale versions offer in principle a way to design realistic immersive virtual worlds that allow multiple users to interact.

REFERENCES

- [1] Tim Brooks, Bill Peebles, Connor Holmes, Will DePue, Yufei Guo, Li Jing, David Schnurr, Joe Taylor, Troy Luhman, Eric Luhman, Clarence Ng, Ricky Wang, and Aditya Ramesh. Video generation models as world simulators. 2024.
- [2] Julie A Dickerson and Bart Kosko. Virtual worlds as fuzzy cognitive maps. *Presence: Teleoperators & Virtual Environments*, 3(2):173–189, 1994.
- [3] Bart Kosko. Fuzzy cognitive maps. *International journal of man-machine studies*, 24(1):65–75, 1986.
- [4] Bart Kosko. Hidden patterns in combined and adaptive knowledge networks. *International Journal of Approximate Reasoning*, 2(4):377–393, 1988.
- [5] 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.
- [6] 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.
- [7] Michael Glykas. *Fuzzy cognitive maps: Advances in theory, methodologies, tools and applications*, volume 247. Springer, 2010.
- [8] 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.
- [9] 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.
- [10] 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.
- [11] A. Adadi and M. Berrada. Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI). *IEEE Access*, 6:52138–52160, 2018.
- [12] Wojciech Samek. *Explainable AI: Interpreting, explaining and visualizing deep learning*, volume 11700. Springer Nature, 2019.
- [13] Akash Kumar Panda, Olaoluwa Adigun, and Bart Kosko. Causal autoencoder-like generation of feedback fuzzy cognitive maps with an LLM agent. In *2025 International Conference on Machine Learning and Applications (ICMLA-2025)*, pages 1234–1241. IEEE, 2025.
- [14] Akash Kumar Panda, Olaoluwa Adigun, and Bart Kosko. The Agentic Leash: Extracting Causal Feedback Fuzzy Cognitive Maps with Mixed Large Language Models. In *International Conference on Large Language Models (CSCI-2025)*. Springer, 2025.
- [15] Akash Kumar Panda, Olaoluwa Adigun, and Bart Kosko. Agentic Chunking and Bayesian De-chunking of AI Generated Fuzzy Cognitive Maps: A Model of the Thucydides Trap. In *IFIP International Conference on Artificial Intelligence Applications and Innovations*, pages 17–32. Springer, 2026.