

The Agentic Leash: Extracting Causal Feedback Fuzzy Cognitive Maps with Mixed Large Language Models

Akash Kumar Panda¹, Olaoluwa Adigun², and Bart Kosko³

¹ University of Southern California, Los Angeles, CA 90007, USA
`akashpan@usc.edu`

² Florida International University, Miami, FL 33199, USA
`olaadigu@fiu.edu`

³ University of Southern California, Los Angeles, CA 90007, USA
`kosko@usc.edu`

Abstract. We design a large-language-model (LLM) agent that extracts causal feedback fuzzy cognitive maps (FCMs) from raw text. The causal learning or extraction process is agentic both because of the LLM’s semi-autonomy and because ultimately the FCM dynamical system’s equilibria drive the LLM agents to fetch and process causal text. The fetched text can in principle modify the adaptive FCM causal structure and so modify the source of its quasi-autonomy—its equilibrium limit cycles and fixed-point attractors. This bidirectional process endows the evolving FCM dynamical system with a degree of autonomy while still staying on its agentic leash. We show in particular that a sequence of three finely tuned system instructions guide an LLM agent as it systematically extracts key nouns and noun phrases from text, as it extracts FCM concept nodes from among those nouns and noun phrases, and then as it extracts or infers partial or fuzzy causal edges between those FCM nodes. We test this FCM generation on a recent essay about the promise of AI from the late diplomat and political theorist Henry Kissinger and his colleagues. This three-step process produced FCM dynamical systems that converged to the same equilibrium limit cycles as did the human-generated FCMs even though the human-generated FCM differed in the number of nodes and edges. A final FCM mixed generated FCMs from separate Gemini and ChatGPT LLM agents. The mixed FCM absorbed the equilibria of its dominant mixture component but also created new equilibria of its own to better approximate the underlying causal dynamical system.

Keywords: fuzzy cognitive maps · causal reasoning · agentic LLMs

1 The Agentic Leash: Growing Fuzzy Cognitive Map Dynamical Systems from Text

We show how agentic passes through structured LLM agents can grow causal feedback fuzzy cognitive maps (FCM) from sampled text documents.

These causal FCM feedback dynamical systems form local fuzzy or partial causal rules from the sampled documents. This local causal structure in turn defines global equilibrium limit cycles that serve as scenario-like answers to causal what-if questions. They also define the very source of the FCM dynamical system’s agency – its evolving equilibrium limit cycles. This differs from ordinary feedforward agentic systems whose agency resides only in programmed commands. Mixing FCMs can give both richer learned causal knowledge bases and richer global equilibria. The extraction process is agentic [1] or partially autonomous because the FCM’s evolving global equilibria command the LLM agents to fetch and process further text that then tends to change the commanding FCM equilibria. Related work uses an autoencoder-like mapping to convert FCMs to text and continue the reverberatory process [13]. This bidirectional process keeps the FCM’s LLM agents on a type of flexible agentic leash.

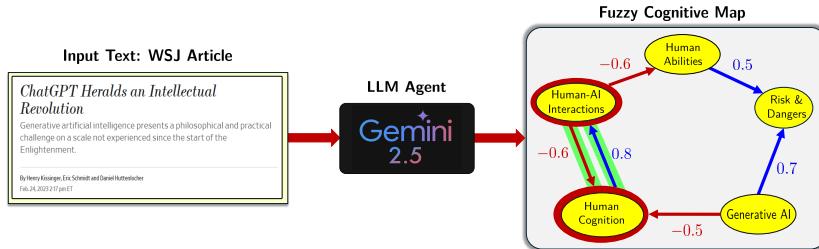


Fig. 1: A Large Language Model (LLM) extracts causal variables and their causal relationships out of a *Wall Street Journal* article from Henry Kissinger and colleagues about the promise of AI and then creates a Fuzzy Cognitive Map (FCM). The figure shows only 5 out of the 15 AI-extracted nodes and the directed weighted edges that connect them. The positive edges are in blue and the negative edges are in red. The figure highlights one of many feedback loops in the FCM in green. In this case: Growth of *Human Cognition* increases *Human-AI Interactions* but an increase in *Human-AI Interactions* decreases *Human Cognition*.

Figure 1 shows a feedback causal sub-network of the complete 15-node FCM in Figure 2. An LLM agent grew the complete FCM from a recent AI article titled “*ChatGPT Heralds an Intellectual Revolution*” by Henry Kissinger et al. in the *Wall Street Journal* [8]. The green highlight shows just one embedded feedback loop that traces the causal flow from *Human-AI Interaction* to *Human Cognition* and then back to itself: *Human-AI Interaction* \rightarrow *Human Cognition* \rightarrow *Human-AI Interaction*. Here “ \rightarrow ” denotes a negative edge and “ \rightarrow ” denotes a positive edge. Even this smaller 5-node sub-FCM encodes equilibrium limit cycles that can serve as answers to what-if questions. Figure 7 below shows this process for a simpler FCM.

These learned FCM knowledge graphs consist of local causal rules. The rules give an immediate local form of interpretability or explainable AI (XAI) [2,4,6,

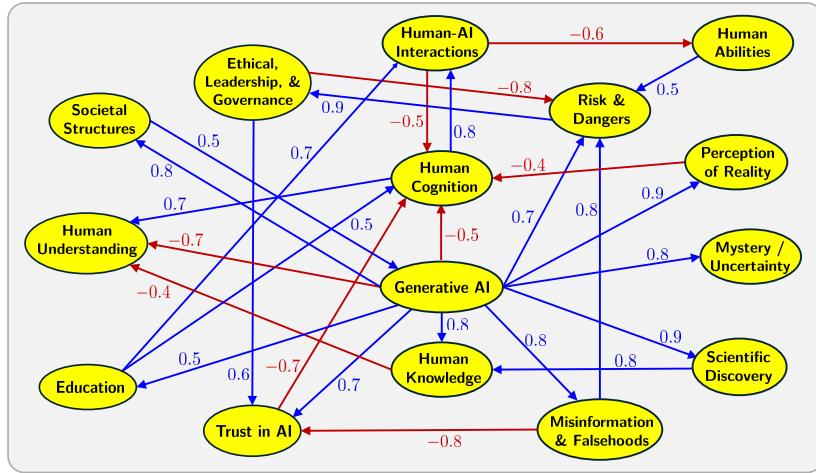


Fig. 2: A 15-node FCM extracted by the LLM from the WSJ article titled “*ChatGPT Herald an Intellectual Revolution*” by Henry Kissinger et al. This FCM converges to a 4-step limit cycle. The limit cycle predicts that the growth of generative AI comes in waves or cycles. In the 1st step generative AI gets widely used, human-AI interactions rise, and human knowledge grows. This helps leaders govern ethically in the 2nd step but it also comes with risks and dangers. In the 3rd step: People trust the mysterious and uncertain AI, misinformation and falsehoods spread, and society changes. But AI improves education at the same time and leads to scientific discoveries. Generative AI is used widely again in the 4th step but this time without ethical leadership. Society changes again but this time without generative AI.

7, 15]. The embedded or hidden limit cycles in feedback FCMs also give a global form of XAI.

A key problem is that text articles usually discuss the causal variables and causal relationships of an underlying dynamical system without explicitly mentioning this system or its dynamics. Authors may have to guess how the feedback causal-rule based system behaves under given conditions. Authors or speakers may not even be aware that their causal portraits define nonlinear feedback dynamical systems. Nor does the author or the reader have an easy way to see the system's dynamics or verify the author's equilibrium predictions even if they are aware of the feedback structure.

A further problem in FCM modeling is that it can be expensive and time consuming to hire an expert to read the text and correctly predict the dynamics of the system. A large-scale FCM may well be too complex for such predictions. Large Language Models (LLMs) can instead parse the body of text and then process it given prompts from users or from the evolving FCM dynamical system itself.

We need an LLM that can systematically extract causal variables and their causal relationships from the text without human supervision. These causal rules from the text should then predict the dynamics of the system described by the text even if the text does not mention the dynamics. FCMs can model this causal information as a weighted directed cyclic graph and also give qualitative

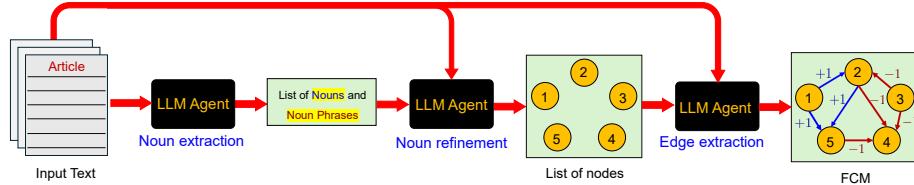


Fig. 3: LLM-based FCM-extraction from text: An LLM agent extracts FCM nodes and edges from a text input through a three-step systematic process. These steps use the same LLM agent with guiding system instructions. The LLM extracts all the nouns and noun phrases used in the text during Step 1. It extracts the FCM nodes from the list of nouns and noun phrases in Step 2. This involves refining the noun list from Step 1 based on associated qualitative or quantitative measures as well as the existence of causal links. The LLM extracts fuzzy-weighted edges from a list of node-pairs to complete the FCM in Step 3.

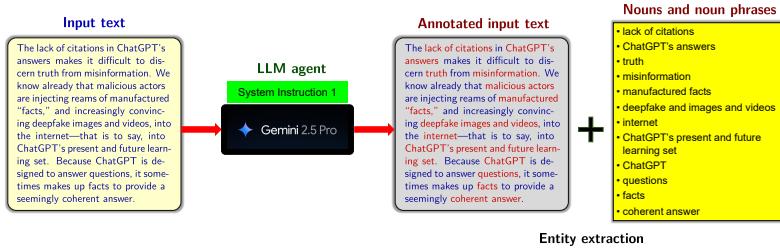


Fig. 4: Noun Extraction with an agentic LLM: The LLM takes a paragraph of text as input and identifies the nouns, noun phrases, and pronouns present in the text. The figure on the right highlights the extracted nouns and noun phrases in red.

equilibrium predictions from the equilibria of the approximating causal FCM [5, 9–11, 11, 12, 14, 16, 17, 19].

We show that LLM agents can systematically extract FCM nodes and edges from a body of text through a sequence of finely tuned system instructions. The LLMs can give reasons for their decisions by quoting from the text source. These textual anchors can also help reduce hallucinations.

Figure 1 shows the simplest version of the FCM agentic set-up. An LLM takes the WSJ article “*ChatGPT Heralds an Intellectual Revolution*” by Henry Kissinger *et alia* as text input and extracts causal variables and causal relationships among those variables from the text. It then builds an FCM out of those causal rules. Figure 1 shows only 5 of the 15 nodes from the extracted FCM. Figure 2 shows the complete FCM.

Section 3 explains FCMs and shows how FCMs can model or approximate causal dynamical systems. The section explains what the FCM nodes and edges mean and how the FCM evolves in discrete time and converges to its equilibria. FCMs can also evolve in continuous time.

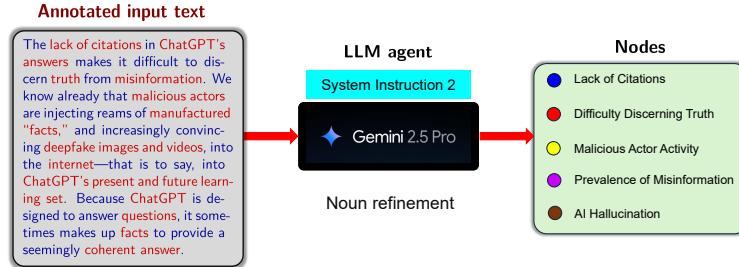


Fig. 5: FCM-node extraction with an agentic LLM: The LLM takes the list of nouns and noun phrases as input and filters those that are associated with some kind of qualitative or quantitative measure. This gives a list of 5 FCM nodes.

Section 4 explains the three-step systematic process to extract an FCM from a body of text that describes a dynamical system. It explains the 3 respective prompts to the LLM agent that detect nouns and noun phrases in the text, that extract nodes from the nouns, and that extract edges from the node-pairs. Figure 3 shows this process.

Section 5 shows examples of FCM extraction. Figures 4-7 show this process for a one-paragraph input text. Figure 8 compares the limit cycles of the extracted FCM to those of a human-generated FCM. The figure shows that both FCMs converge to the same limit cycle. Figure 2 shows the complete FCM extracted from Henry Kissinger’s WSJ article. Figure 12 shows the corresponding limit cycle that the author implied despite never explicitly mentioning it in the article.

Figure 14 mixes the extracted-FCMs from the same article with the same guided prompts but used 2 different LLM agents: Gemini-2.5 Pro and ChatGPT-4.1. The 24-node mixed FCM used equal mixing weights for both FCMs. Figure 13 mixes the corresponding edge matrices of the 2 mixture-component FCMs using zero-padding and convex combination. Figure 15 shows the limit cycles of this final mixed FCM.

2 Extracting FCMs from Text and Related Work

This paper shows how to extract FCM nodes and edges from raw text by just using the language of the text. It does not rely on labels or annotations. The FCM nodes are usually nouns or noun phrases associated with some qualitative or quantitative measure. Verbs in the text describe FCM causal edges. Edge extraction is transparent and interpretable because it gives direct evidence for the edges by quoting the text.

Prior work by Maryam Berijanian et al. [3] also extracted FCMs from text sources using LLMs. But they trained their LLM through supervision and annotated their text input with (“source”, “target”, “direction”) tuples. This may help the LLMs extract FCMs but it defeats the whole purpose of *automated* FCM extraction because it requires a human to read the text and annotate the nodes.

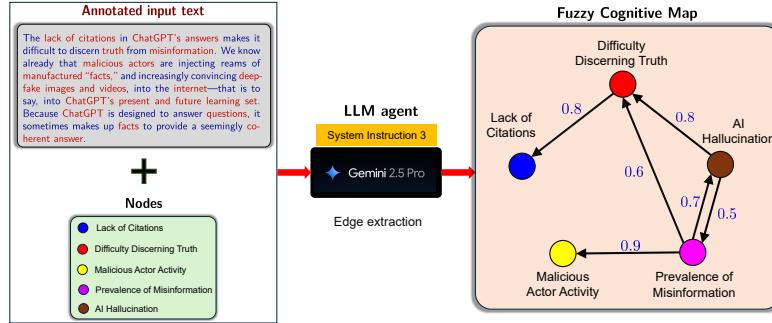


Fig. 6: FCM-edge extraction with an agentic LLM: The LLM goes through node-pairs and looks for textual evidence for causal connection. This gives the FCM 6 causal edges.

Extracting FCM nodes from text somewhat resembles Named Entity Recognition (NER) tasks for LLMs [18]. Here the LLMs identify key information from the text and classify them based on predefined categories. The LLMs can extract nouns and noun phrases from a text based on the properties of nouns in that language. But every noun or noun phrase in the text does not correspond to a node in the FCM. That approach would grow impossibly large and inaccurate FCMs. Finely-tuned system instructions can prompt the LLM to pick only the nouns that also share the properties of an FCM node.

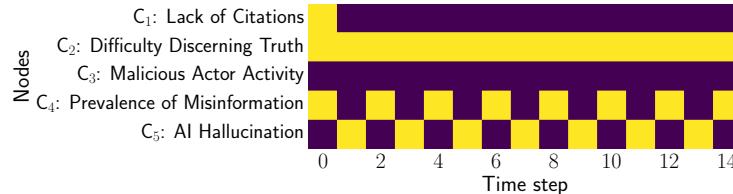
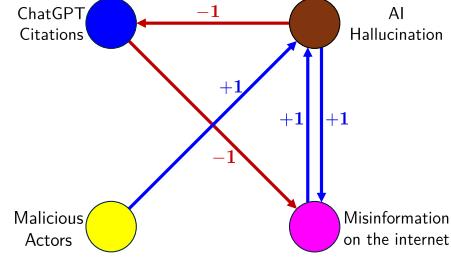


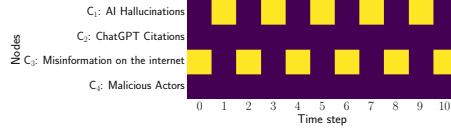
Fig. 7: The limit cycles of LLM's extracted FCM. The LLM's FCM converged to a 2-step limit cycle starting from the initial state $[1 \ 1 \ 0 \ 1 \ 0]$. The nodes are along the x -axis and the time steps are along the y -axis. Active nodes are in yellow and the inactive nodes are in purple.

Unguided prompts into an LLM are not reliable for extracting FCMs in general. They often have trouble distinguishing negative causal relationships from positive ones and might well inaccurately weight the FCM edges. They also often do not give reasons for the FCM edges. They may even hallucinate edges based on things not explicitly mentioned in the text itself.

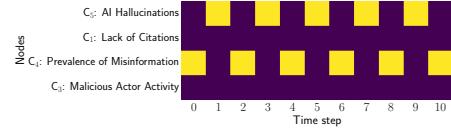
We design system instructions that force the LLMs to choose the nodes and edges in a systematic manner. This transparent 3-step process explains all the decisions made by the LLM and also reduces the chances and extent of hallucinations. The LLM supports all its decisions by quoting evidence from the text source.



(a) Human-generated FCM



(b) Limit cycle: Human-generated FCM



(c) Limit cycle: LLM-generated FCM

Fig. 8: FCM extraction from a Kissinger AI-hallucination text: (a) Human-generated FCM from the same “AI hallucination” input text. (b) Limit cycles from the human-generated FCM: The time steps are along the x -axis and the nodes are along the y -axis. Active nodes are in yellow and the inactive nodes are in purple. (c) Limit cycles from Gemini’s extracted FCM for similar initial conditions. The figure only shows 4 nodes from Gemini’s FCM that correspond to the nodes from human-generated FCM.

3 FCM Basics

FCMs model causal dynamical systems as directed weighted graphs. 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.

3.1 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 says “ C_i causes C_j ”. The fuzzy edge-weights describe partial causality. The edge weight $w_{ij} \in [-1, 1]$ on the edge e_{ij} gives the degree to which C_i causes C_j :

$$w_{ij} = \text{degree}(C_i \rightarrow C_j). \quad (1)$$

A positive w_{ij} says that C_j increases when C_i increases. A negative w_{ij} means that C_j decreases if C_i increases. The weight w_{ij} is zero when there is no causal edge between C_i and C_j . The magnitude of w_{ij} is high when there is a strong causal relationship between C_i and C_j . Low magnitude of w_{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 edge weight w_{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.

3.2 Discrete Time Evolution of FCMs

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” 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” 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.

FCMs evolve in discrete time through vector-matrix multiplication and nonlinear squashing. The state $C_j(t+1)$ of the j^{th} concept node C_j at discrete time step $t+1$ is:

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

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

The sum $\sum_{i=1}^n C_i(t) w_{ij}$ is the matrix product between the state row-vector $C(t)$ and the edge matrix E . The nonlinear function Φ then squashes this product between zero and one.

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.

Enlightenment science accumulated certainties; the new AI generates cumulative ambiguities. Enlightenment science evolved by making mysteries explicable, delineating the boundaries of human knowledge and understanding as they moved. The two faculties moved in tandem: Hypothesis was understanding ready to become knowledge; induction was knowledge turning into understanding. In this Age of AI, riddles are solved by processes that remain unknown. This disorienting paradox makes mysteries unmysterious but also unexplainable. Inherently, highly complex AI furthers human knowledge but not human understanding—a phenomenon contrary to almost all of post-enlightenment modernity.

Fig. 9: A paragraph from the WSJ article *ChatGPT Heralds an Intellectual Revolution* by Henry Kissinger et al.

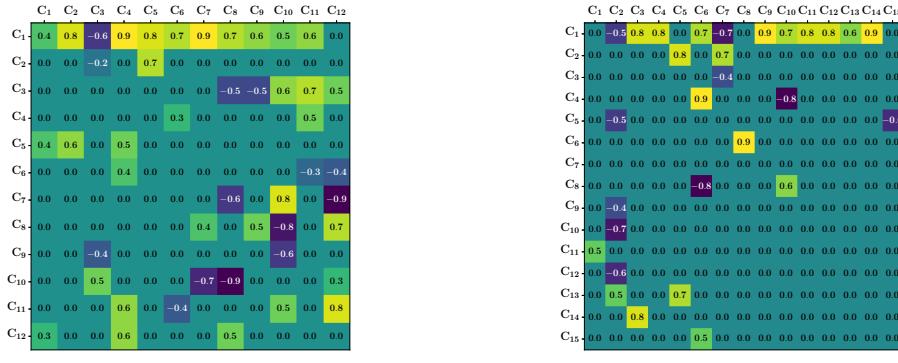


Fig. 10: FCM with Gemini-2.5 Pro: (a) The edge matrix corresponding to the 12-node Kissinger AI FCM from an unguided prompt. (b) The edge matrix corresponding to the 15-node Kissinger AI FCM from a guided prompt.

3.3 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.

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.

3.4 FCM Mixtures

FCMs combine their knowledge through convex mixing. Consider m FCMs. Say the k^{th} FCM has the set of nodes S_k and the edge matrix E_k . The node-set S

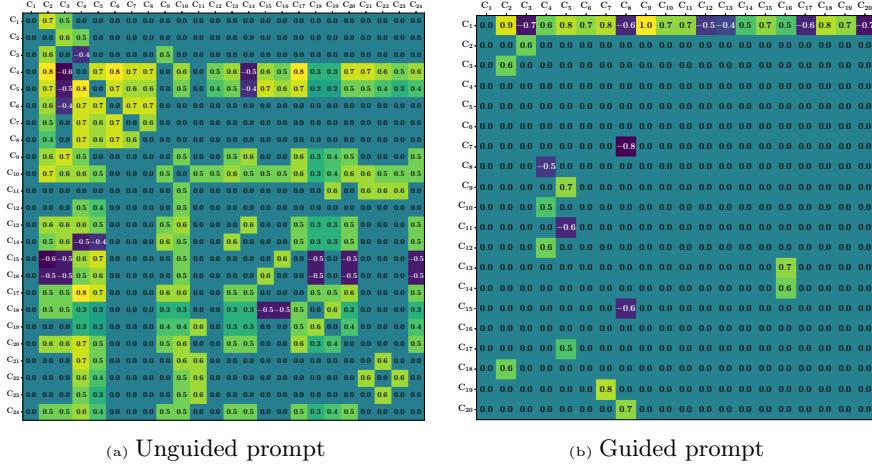


Fig. 11: FCM with GPT-4.1: (a) The edge matrix corresponding to the 24-node Kissinger AI FCM with an unguided prompt . (b) The edge matrix corresponding to the 20-node Kissinger AI FCM with a guided prompt.

for the N -node FCM mixture is $S_1 \cup S_2 \cup S_3 \cup \dots \cup S_m$. The $N \times N$ matrix \tilde{E}_k pads the k^{th} edge matrix E_k with zero rows and zero columns corresponding to the nodes in the set difference $S - S_k$. The edge matrix E for the mixed FCM is

$$E = \sum_{k=1}^m v_k \tilde{E}_k \quad (3)$$

where v_k are convex mixing weights such that $v_k \geq 0$ and $\sum_{k=1}^m v_k = 1$. FCM mixing is closed. This means that mixing FCMs gives back an FCM. Figure 13 shows this process for a 2-component equal-weighted FCM-mixture.

4 LLM-Based FCM Extractor

Text articles often describe causal dynamical systems without explicit mathematical formulation. This makes it hard to understand what dynamics the underlying dynamical system predicts. It is also difficult to verify the author’s guess regarding the dynamical system’s behavior. A human expert can read through the text and model the dynamical system it describes. But this may not always be possible, may take too long, or may be too expensive. This motivates automated dynamical-system modeling from text.

LLMs are good at parsing text and processing natural language. They can extract causal information from text that describes dynamical systems and use it to model the underlying dynamical systems as FCMs. But LLMs hallucinate with unguided prompts. Their answer may also vary across multiple runs with the same prompt. The LLMs need guiding prompts that are designed to extract FCMs.

System instructions manipulate an LLM agent to behave a certain way. System instructions tell the LLM how to process the input and how to structure the output. The same LLM can react differently to “prompts” based on different system instructions. A sequence of appropriate system instructions can guide an LLM agent to systematically extract FCM-nodes and edges from text input.

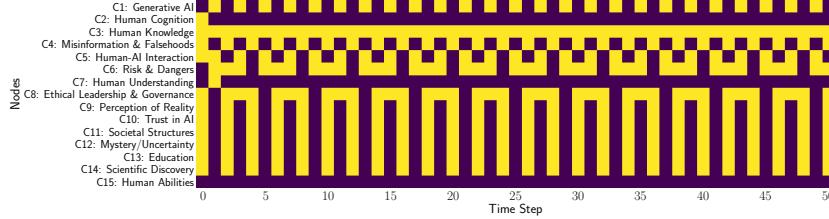


Fig. 12: Limit cycles from the extracted FCM. The time steps are along the x -axis and the nodes are along the y -axis. Active nodes are in yellow and the inactive nodes are in purple.

4.1 Step 1: Noun Extraction

The nodes of an FCM are present in the text as nouns and noun phrases. System instructions ask the LLM Agent to go through the text sentence-by-sentence and make a list of all nouns, noun phrases and pronouns present in the text. The agent then replaces all the pronouns with their corresponding antecedents as well.

4.2 Step 2: Noun Refinement

FCM nodes are associated with a qualitative or quantitative measure. This measure can causally increase or decrease due to other nodes. The 2nd system instruction to the LLM agent asks it to filter the list of nouns and noun phrases and only keep nouns that are associated with qualitative or quantitative measures and are causally connected to other nouns with similar measures. It then returns a refined list of nouns and noun phrases that will serve as the nodes of the FCM.

4.3 Step 3: Edge Extraction

The LLM agent with the 3rd set of system instructions pairs every node from the refined list with every other node. It then goes through every node-pair and looks for verbs from the text to establish a positive, negative, or zero causal connection. The agent also gives reason for each edge by directly quoting from the text. It also assigns a weight to the edge based on the verb used in the text. This gives us the complete FCM described by the text.

Equation (3) then gives the limit cycles of the FCM. These are the dynamical equilibria that the text suggests.

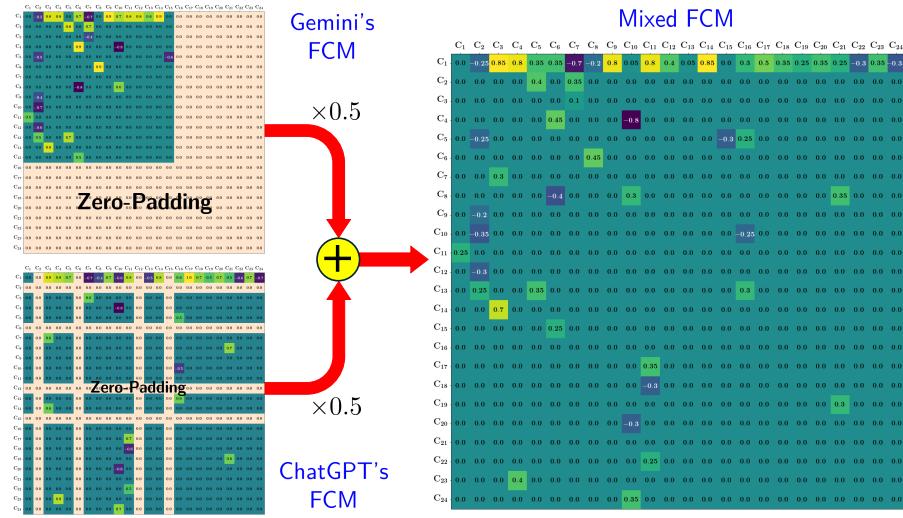


Fig. 13: FCMs' edge matrices mix through zero-padding and convex combination described in equation (3). The edge matrix of the 15-node Gemini-extracted FCM is on the top-left and the edge matrix of the 20-node ChatGPT-extracted FCM is on the bottom-left. The edge matrix of the 24-node equal-weight FCM-mixture is on the right. The bigger edge weights are colored brighter in the matrices. The matrices on the left also highlight the zero-padded rows and columns of the mixture components.

5 Simulations

Gemini-2.5 Pro LLM with “temperature = 1” and “top $p = 0.95$ ” took a paragraph of text from the Kissinger article describing the effects of misinformation on the internet on AI hallucinations as input. It then followed system instructions 1-3 respectively to systematically extract the nouns, nodes, and causal edges. This gave a FCM with 5 nodes and 6 edges.

The LLM highlighted the nouns and noun phrases directly from the text as reasons for choosing certain nodes. Figures 4 and 5 highlight these nouns and noun phrases in red. The node “*Malicious Actor Activity*” for example comes from the noun phrase “*malicious actors*”. The negative edge from “*lack of citations*” to “*Difficulty discerning truth*” comes from the sentence “*The lack of citations in ChatGPT's answers makes it difficult to discern truth from misinformation*”. Figures 4-6 show this 3-step process step-by-step.

This FCM converged to fixed points or limit cycles depending on the initial state. AI hallucinations spread misinformation over the internet. AI systems then train on this misinformation which causes them to hallucinate. Both steps make truth harder to discern. Figure 7 also shows a 2-step limit cycle of the extracted FCM.

A human took the same text as input but came up with an FCM with only 4 nodes and 5 edges. The 4 human-generated nodes “*AI Hallucinations*”, “*Misinformation on the internet*”, “*ChatGPT Citations*”, and “*Malicious Actors*”

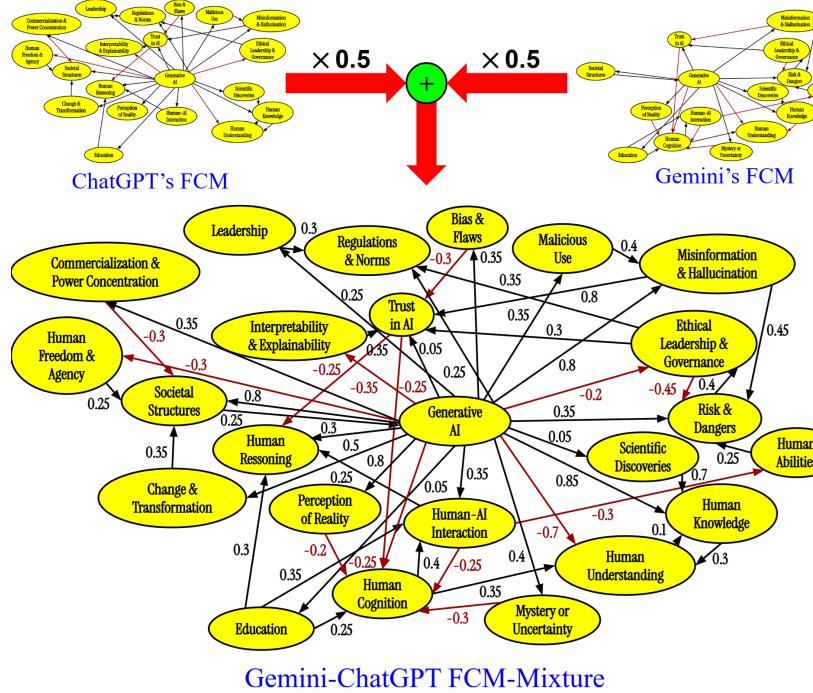


Fig. 14: Mixed FCMs: An equal-weight mixture of the 15-node Gemini-extracted FCM and the 20-node ChatGPT-extracted FCM after appropriate zero-padding of rows and columns. The final mixed FCM has 24 nodes and 52 edges. The 15-node Gemini-extracted FCM is on the top-right and the 20-node ChatGPT-extracted FCM is on the top-left.

are also present in the LLM-extracted FCM. The node “*ChatGPT Citations*” plays the opposite role of “*Lack of Citations*”. This human-generated FCM also converged to a 2-step limit cycle. Figure 8 compares this limit cycle to that of the LLM-extracted FCM by only considering the nodes that are present in both the FCMs. It shows that both FCMs converge to the same limit cycle.

The same LLM then took a WSJ article “*ChatGPT Heralds an Intellectual Revolution*” by Henry Kissinger et al. as input. Figure 9 shows one paragraph from this article. An unguided-prompt to the LLM extracted a FCM with 12 nodes. Figure 10a shows the edge matrix corresponding to this FCM. A guided prompt to the same LLM gave a 15-node FCM instead. Figure 2 shows this FCM and Figure 10b gives its edge matrix. This 15-node FCM converged to a 4-step limit cycle. Figure 12 shows this limit cycle.

This limit cycle also tells a what-if story of the underlying dynamical system. Generative AI gets widely used, human-AI interactions rise, and human knowledge grows. This helps leaders govern ethically but it also comes with risks and dangers. People trust the mysterious and uncertain AI, misinformation and falsehoods spread, and society changes. But AI improves education at the same time leading to scientific discoveries. Generative AI is used widely again but this

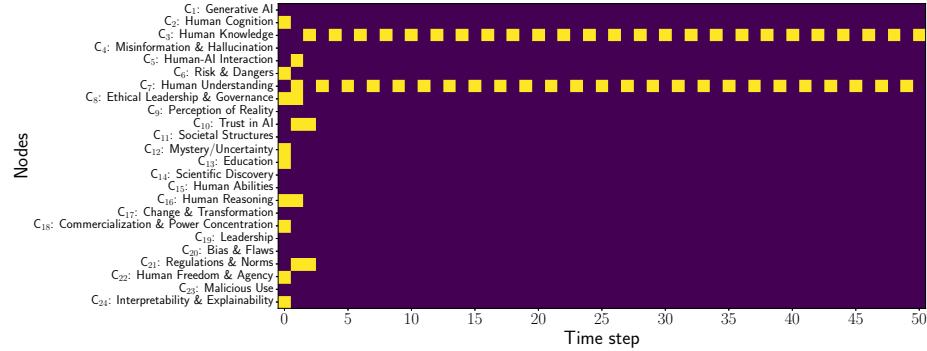


Fig. 15: Limit cycle from the mixed Gemini-ChatGPT FCM. The time steps are along the x -axis and the nodes are along the y -axis. Active nodes are in yellow and the inactive nodes are in purple. The 2-step limit cycle oscillates between states “Human Knowledge” and “Human Understanding.”

time without ethical leadership. Society changes again but this time without generative AI. This limit cycle was not mentioned in the source article at all but the article did imply it through its discussion on the causal variables and their relationships.

OpenAI’s ChatGPT-4.1 LLM agent respectively produced a 24-node FCM and a 20-node FCM based on the same unguided and guided prompts. Figure 11 shows the edge matrices corresponding to these FCMs. The 15-node Gemini-FCM and the 20-node ChatGPT-FCM mixed with equal mixing weights to give the final 24-node mixed FCM. Figure 14 shows this FCM-mixture. Figure 13 shows its edge matrix while Figure 15 shows its limit cycle.

6 Conclusions

This paper shows how iterated and structured LLM agents can generate feedback causal fuzzy cognitive maps from text documents or from transcribed speech. Further causal structuring should produce larger FCMs and more powerful LLMs should produce richer and more accurate FCMs. The process scales across text sources and knowledge bases because causal language is fairly standard and because FCM operation depends only vector-matrix multiplication and simple nonlinear units.

The mixing structure of FCMs also allows a given large text to break into chunks and subchunks with corresponding mixed FCM components. This process can help with node and edge identification in large documents. Mixing such FCMs still produces an FCM and this should further help the growth and uses of very-large-scale causal knowledge networks. More advanced agentic FCM systems will lengthen the agentic leash and allow the systems to grow and use complex causal networks as it both shapes and obeys the evolving global equilibria of the total FCM dynamical system.

References

1. Acharya, D.B., Kuppan, K., Divya, B.: Agentic AI: Autonomous intelligence for complex goals—a comprehensive survey. *IEEE Access* (2025)
2. Adadi, A., Berrada, M.: Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI). *IEEE Access* **6**, 52138–52160 (2018)
3. Berijanian, M., Dork, S., Singh, K., Millikan, M.R., Riggs, A., Swaminathan, A., Gibbs, S.L., Friedman, S.E., Brugnone, N.: Soft measures for extracting causal collective intelligence. *arXiv preprint arXiv:2409.18911* (2024)
4. Giammattei, L., Guerriero, A., Pietrantuono, R., Russo, S.: Causal reasoning in software quality assurance: A systematic review. *Information and Software Technology* **178**, 107599 (2025)
5. Glykas, M.: Fuzzy cognitive maps: Advances in theory, methodologies, tools and applications, vol. 247. Springer (2010)
6. Guimerà, R., Reichardt, I., Aguilar-Mogas, A., Massucci, F.A., Miranda, M., Palàrè, J., Sales-Pardo, M.: A Bayesian machine scientist to aid in the solution of challenging scientific problems. *Science Advances* **6**(5) (2020)
7. Horst, F., Lapuschkin, S., Samek, W., Müller, K.R., Schöllhorn, W.I.: Explaining the unique nature of individual gait patterns with deep learning. *Scientific Reports* **9**(1), 1–13 (2019)
8. Kissinger, H., Schmidt, E., Huttenlocher, D.: Chatgpt heralds an intellectual revolution (2023)
9. Kosko, B.: Differential Hebbian learning. In: AIP Conference proceedings. vol. 151, pp. 277–282. American Institute of Physics (1986)
10. Kosko, B.: Fuzzy cognitive maps. *International journal of man-machine studies* **24**(1), 65–75 (1986)
11. Kosko, B.: Hidden patterns in combined and adaptive knowledge networks. *International Journal of Approximate Reasoning* **2**(4), 377–393 (1988)
12. Osoba, O.A., Kosko, B.: Fuzzy cognitive maps of public support for insurgency and terrorism. *The Journal of Defense Modeling and Simulation* **14**(1), 17–32 (2017)
13. Panda, A.K., Adigun, O., Kosko, B.: Causal autoencoder-like generation of feedback fuzzy cognitive maps with an llm agent. In: 2025 24th IEEE International Conference on Machine Learning and Applications (ICMLA). pp. 1234–1241. IEEE (2025)
14. Papageorgiou, E.I.: Fuzzy cognitive maps for applied sciences and engineering: from fundamentals to extensions and learning algorithms, vol. 54. Springer Science & Business Media (2013)
15. Samek, W.: Explainable AI: interpreting, explaining and visualizing deep learning, vol. 11700. Springer Nature (2019)
16. Stach, W., Kurgan, L., Pedrycz, W.: A divide and conquer method for learning large fuzzy cognitive maps. *Fuzzy Sets and Systems* **161**(19), 2515–2532 (2010)
17. Taber, R., Yager, R.R., Helgason, C.M.: Quantization effects on the equilibrium behavior of combined fuzzy cognitive maps. *International Journal of Intelligent Systems* **22**(2), 181–202 (2007)
18. Wang, S., Sun, X., Li, X., Ouyang, R., Wu, F., Zhang, T., Li, J., Wang, G.: Gpt-ner: Named entity recognition via large language models. *arXiv preprint arXiv:2304.10428* (2023)
19. Ziv, G., Watson, E., Young, D., Howard, D.C., Larcom, S.T., Tanentzap, A.J.: 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)