

Agentic Chunking and Bayesian De-chunking of AI Generated Fuzzy Cognitive Maps: A Model of the Thucydides Trap

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 automatically generate feedback causal fuzzy cognitive maps (FCMs) from text by teaching large-language-model agents to break the text into overlapping chunks of text. Convex mixing of these chunk FCMs gives a representative cyclic FCM knowledge graph. The text chunks can have different levels of overlap. The chunk FCMs still mix to form a new FCM causal knowledge graph. The mixing technique scales because it uses light computation with sparse causal chunk matrices. The mixing structure allows an operator-level type of Bayesian inference that produces “de-chunked” or posterior-like FCMs from the mixed FCM. These de-chunked FCMs are useful in their own right and allow further iterations of Bayesian updating. We demonstrate these mixing techniques on the essay text of Allison’s “Thucydides Trap” model of conflict between a dominant power such as the United States and a rising power such as China. The FCM dynamical systems predict outcomes as they equilibrate to fixed-point or limit-cycle attractors. Seven out of 8 FCM knowledge graphs predicted a type of war when we stimulated them by turning on and keeping on the concept node that stands for the rising power’s ambition and entitlement. Gemini 3.1 LLMs served as the chunking AI agents.

Keywords: Agentic AI · fuzzy cognitive maps · text chunking · LLM Agent · Bayesian posterior FCMs.

1 AI Chunking and Bayesian De-chunking of FCMs

We show how to use large-language-model agents to break sampled text into overlapping chunks and then map those text chunks to sparse feedback fuzzy cognitive maps (FCMs). The technique offers a practical way to decompose large text documents into representative knowledge graphs and then recombine them into a unified graph that represents the entire document. Users can then ask this combined FCM network what-if causal questions and get equilibrium answers.




Fig. 1: AI-automated chunking and de-chunking of sampled text to create chunk-based Fuzzy Cognitive Maps (FCMs) and the Bayesian-like posterior FCMs. The chunking experiment uses Graham Allison’s article in *The Atlantic* titled “*The Thucydides Trap: Are the U.S. and China headed for war?*” as input. Allison explains: “*When a rising power is threatening to displace a ruling power, standard crises that would otherwise be contained, like the assassination of an archduke in 1914, can initiate a cascade of reactions that, in turn, produce outcomes none of the parties would otherwise have chosen.*” The chunking process splits Allison’s article into n overlapping chunks of contiguous text. Then the large language model (LLM) Gemini 3.1 extracts the respective FCMs $\mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_n$ from these text chunks using the customized system instructions in Section 3. These n FCMs then mix through convex combination to give the mixed FCM $\mathcal{F}_{\text{Likelihood}}$. The de-chunking process finds the posterior FCM $\mathcal{P}_k$ that corresponds to the likelihood FCM $\mathcal{F}_k$. Then the posterior FCMs also mix through convex combination to give $\mathcal{F}_{\text{Posterior}}$.

The FCMs are fuzzy because their local directed causal edges take values in the bipolar interval $[-1, 1]$ and represent partial causality or degrees of causal decrease or increase. The FCMs are cyclic knowledge graphs that define nonlinear feedback dynamical systems whose equilibria give fixed-point or limit-cycle answers to the what-if questions that users impose on their causal concept nodes.

Figure 1 shows how text from Allison’s “Thucydides Trap” essay [1] on Great-Power conflict leads to 3 chunk FCMs that combine into a representative FCM for the entire sampled document. Figure 2 shows how these chunk FCMs combine to form a representative FCM knowledge graph by mixing their zero-padded edge matrices with convex mixing weights or probabilities. This automated and mixed Thucydides FCM differs in kind from the earlier hand-drawn FCM in [8].

We further show that we can de-mix or “de-chunk” this mixed FCM to give Bayesian-like (operator-based) posterior FCMs as in Figure 3. These FCMs can then participate in future bouts of Bayesian-like updating and mixing.

The next sections explain how FCM dynamical systems work and how to guide LLM agents to perform semantic-split-like text chunking and the crucial mapping from text chunks to combinable FCMs. We demonstrate these techniques on Allison’s popular essay that addresses whether the United States and China may get caught in a warlike Thucydides Trap.




Fig. 2: Likelihood FCM Mixture. The 3 likelihood FCMs $\mathcal{F}_1$, $\mathcal{F}_2$, and $\mathcal{F}_3$ extracted from chunks of the “Thucydides Trap” article are on the top. The 11-node mixture FCM $\mathcal{F}_{\text{Likelihood}}$ mixes the 3 likelihood FCMs with equal mixing weights $\frac{1}{3}$.

2 Fuzzy Cognitive Maps

FCMs model causal dynamical systems as directed weighted graphs [3, 5–7, 9–15]. 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.

In contrast: Bayesian belief networks do not mix in general. They are directed-*acyclic* graphs that lack feedback and so have no dynamics. Inference requires marginalization of joint probabilities that involves NP-hard computation [2].

2.1 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 means “ C_i causes C_j ”. The fuzzy edge-weights describe partial causality. The edge weight $e_{ij} \in [-1, 1]$ on the edge from the i^{th} node to the j^{th} node gives the degree to which C_i causes C_j : $e_{ij} = \text{Degree}(C_i \rightarrow C_j)$.

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 . Low magnitude of e_{ij} describes a weak causal relationship between C_i and C_j .




Fig. 3: FCM De-chunking and Bayesian-like posterior FCMs. Equation (3) converts the mixed likelihood FCM edge matrix E^l into 3 sparse posterior-like edge matrices E_1^p , E_2^p , and E_3^p . These posterior matrices then define the posterior FCMs $\mathcal{P}_1$, $\mathcal{P}_2$, and $\mathcal{P}_3$.

An $n \times n$ matrix E describes all the directed weighted edges of a n -node FCM. The 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.

2.2 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.

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:

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

where 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 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.

2.3 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.

2.4 Clamping and 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.

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 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.

2.5 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 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

$$E = \sum_{k=1}^m w_k \tilde{E}_k \quad (2)$$

where w_k are convex mixing weights such that $w_k \geq 0$ and $\sum_{k=1}^m w_k = 1$. FCM mixing is closed: Mixing FCMs always gives back an FCM.

2.6 Posterior FCMs from the Ratio Structure of Bayes Theorem

Total probability mixes likelihoods $P(E|H_j)$ using the prior probabilities $P(H_j)$ as convex mixing weights: $P(E) = \sum_{j=1}^m P(H_j)P(E|H_j)$ for evidence E and disjoint exhaustive hypotheses H_j . Then Bayes theorem gives the posterior probabilities $P(H_k|E)$: $P(H_k|E) = P(H_k)P(E|H_k)/P(E)$ that further simplifies to $P(H_k|E) = P(H_k)P(E|H_k)/\sum_{j=1}^m P(H_j)P(E|H_j)$ because of total probability. The Bayes posterior $P(H_k|E)$ has the ratio form $\frac{a}{a+b} = a(a+b)^{-1}$ and this we can generalize with FCM nonlinear operators and their causal edge matrices.

Equation (2) shows that mixed FCMs also share the same convex mixing structure as does total probability. But they mix edge *matrices* instead of mixing likelihood probabilities. This gives a new Bayesian-like posterior FCM. Say m likelihood FCMs $F_1, F_2, \dots, F_m$ with edge matrices $E_1^l, E_2^l, \dots, E_m^l$ mix with respective mixing weights $w_1, w_2, \dots, w_m$. Then the posterior FCM P_k that corresponds to the k^{th} FCM F_k has the edge matrix E_k^p :

$$E_k^p = w_k \tilde{E}_k^l (E^l)^{-1} = w_k \tilde{E}_k^l \left(\sum_{j=1}^m w_j \tilde{E}_j^l \right)^{-1} \quad (3)$$

if the mixed-FCM’s edge matrix $E^l = \sum_{j=1}^m w_j \tilde{E}_j^l$ is invertible. So the m posterior edge matrices E_k^p sum to the identity matrix I just as the m posteriors

$P(H_k|E)$ sum to 1: $\sum_{k=1}^m E_k^p = \mathbf{I}$. The posterior FCMs also mix to give the posterior-FCM mixture $\mathcal{F}_{\text{Likelihood}}$ that gives an operator form of total probability: $E^p = \sum_{k=1}^m v_k E_k^p$ for convex mixing weights: $v_k \geq 0$ and $\sum_{k=1}^m v_k = 1$.

If E^l is not invertible then we use the *pseudo-inverse* of E^l :

$$E_k^p = w_k \tilde{E}_k^l (E^l)^T (E^l (E^l)^T)^{-1} \quad (4)$$

$$= w_k \tilde{E}_k^l \left(\sum_{j=1}^m w_j \tilde{E}_j^l \right)^T \left(\left(\sum_{j=1}^m w_j \tilde{E}_j^l \right) \left(\sum_{j=1}^m w_j \tilde{E}_j^l \right)^T \right)^{-1} \quad (5)$$

The posterior FCM may have a few edge weights outside the bipolar interval $[-1, 1]$. We can set their magnitude equal to 1 or renormalize.

3 Mapping Text Chunks to FCMs

This section presents the multi-step process that converts text to FCMs. The process breaks a full text into small chunks and uses an LLM agent to extract the nodes and fuzzy edge weights for each corresponding FCM.




Fig. 4: Binary recursive text splitting: The binary recursive text method breaks down an input text $\mathcal{X}$ to small, contiguous, and non-overlapping text chunks. This method uses a binary tree structure and the text chunks are the leaf nodes at depth $d_{\max}$. The maximum depth $d_{\max} = 3$ so the number of text chunks is $2^3 = 8$.

3.1 Text Chunking

Text chunking is the process of breaking a text down into small, non-overlapping or overlapping text chunks. We define two types of text chunking: non-overlapping chunking and overlapping chunking. Both techniques use a simple binary recursive text splitting method to generate contiguous and small text chunks. Non-overlapping chunking does not generate any overlapping chunks while overlapping chunking does.

Binary Recursive Text Splitting: This method uses a binary tree to recursively divide a text based on the number of paragraphs in small text chunks. The division continues until it reaches the given user-specified maximum depth $d_{\max}$. The leaves of the binary tree are the small and non-overlapping text chunks of the input text. The union or concatenation of these chunks gives back the text at any depth of the binary tree.

The binary text splitting uses the SPLIT function to recursively break down a given input text $\mathcal{X}$ with N paragraphs and a maximum depth $d_{\max}$. The range of $\mathcal{X}$ is $[\text{lo}, \text{hi}]$ with $\text{lo} = 1$ and $\text{hi} = N$. This is the initial range for the splitting and the corresponding depth is $d = 0$. The SPLIT function bisects the current range at the midpoint: $\text{mid} = \lfloor (\text{lo} + \text{hi})/2 \rfloor$. The method then applies SPLIT to these two ranges: $[\text{lo}, \text{mid}]$ and $[\text{mid} + 1, \text{hi}]$, and updates the d . Each node of receives the index $\mathcal{C}(d, k)$ where d is the depth and k is its left-to-right position. This produces $2^{d_{\max}}$ non-overlapping chunks $\mathcal{C}(d_{\max}, 1), \dots, \mathcal{C}(d_{\max}, 2^{d_{\max}})$ that partitions $\mathcal{X}$. Figure 4 shows an example of how this splitting method works.

Non-overlapping Text Chunking: The non-overlapping chunking uses the binary recursive splitting method to divide $\mathcal{X}$ into $2^{d_{\max}}$ small and non-overlapping chunks. Algorithm 1 shows the pseudo-code for this chunking method. It runs recursively and returns the non-overlapping chunks or their corresponding paragraph ranges at the leaf nodes.

Algorithm 1 Non-overlapping Chunking

Require: Text $\mathcal{X}$ with N paragraphs, maximum depth $d_{\max}$

Ensure: Set of non-overlapping chunks $\{\mathcal{C}(d_{\max}, k) : k = 1, \dots, 2^{d_{\max}}\}$

```

1: SPLIT(1, N, 1, 1) ▷ Recursive binary splitting

2: function SPLIT(lo, hi, d, k)
3:   if  $d = d_{\max}$  then
4:      $\mathcal{C}(d_{\max}, k) \leftarrow \mathcal{X}[\text{lo} : \text{hi}]$ 
5:     return
6:   end if
7:    $\text{mid} \leftarrow \lfloor (\text{lo} + \text{hi})/2 \rfloor$ 
8:   SPLIT(lo, mid, d + 1, 2k - 1)
9:   SPLIT(mid + 1, hi, d + 1, 2k)
10: end function

```

The size $N(d, k)$ of chunk $\mathcal{C}(d, k)$ is the number of paragraphs in the chunk. The text chunks at any depth d are such that:

$$\sum_{k=1}^{2^d} |\mathcal{C}(d, k)| = \sum_{k=1}^{2^d} N(d, k) = N \quad (6)$$

where $k, l \in \{1, 2, \dots, 2^d\}$.

Overlapping Text Chunking: This method breaks down an input text to overlapping text chunks. These are the inputs for this method: input text $\mathcal{X}$,

maximum depth $d_{\max}$, and split factors α_1 and α_2 . The method returns $2 \cdot 2^{d_{\max}} - 1$ text chunks in total. Algorithm 2 shows the pseudo-code for this method.

This chunk methods involves two steps. The first step uses the binary recursive text splitting method to break $\mathcal{X}$ to $2^{d_{\max}}$ non-overlapping chunks. These are the leaf nodes for the corresponding binary tree. The second step then creates $2^{d_{\max}} - 1$ overlapping chunks from the non-overlapping chunks. The degree of overlap for the overlapping chunks depends on the values of α_1 and α_2 .




Fig. 5: Overlapping text chunking: This illustrates the overlapping chunking method (Algorithm 2) with $d_{\max} = 3$. This breaks down an input text $\mathcal{X}$ to a sequence of overlapping chunks. The first step applies the binary recursive text split to generate the non-overlapping chunks: $\mathcal{C}(3, 1), \dots, \mathcal{C}(3, 8)$. The second step extracts proportional overlapping chunks $\mathcal{O}_k(\alpha)$ for the non-overlapping chunks.

Each text chunk breaks down to two disjoint sub-chunks: a leading sub-chunk and a trailing sub-chunk. The size of this sub-chunks depend on the split factor α . The leading sub-chunk $\mathcal{L}(\mathcal{C}(d, k), \alpha)$ of $\mathcal{C}(d, k)$ with split factor α is the first $\lceil \alpha \cdot N(d, k) \rceil$ paragraphs of $\mathcal{C}(d, k)$. We have $\ell(d, k) = \lceil \alpha \cdot N(d, k) \rceil$ and $\mathcal{L}(\mathcal{C}(d, k), \alpha) = \mathcal{C}(d, k)[1 : \ell(d, k)]$.

The trailing sub-chunk $\mathcal{T}(\mathcal{C}(d, k), \alpha)$ of $\mathcal{C}(d, k)$ with the same split factor is the last $\lfloor (1 - \alpha) \cdot N(d, k) \rfloor$ paragraphs of $\mathcal{C}(d, k)$. We have $\tau(d, k) = \lfloor (1 - \alpha) \cdot N(d, k) \rfloor$ and $\mathcal{T}(\mathcal{C}(d, k), \alpha) = \mathcal{C}(d, k)[\tau(d, k) : N(d, k)]$.

The overlapping text chunking method generates overlapping text chunks from the non-overlapping text chunks. The first step yields $2^{d_{\max}}$ non-overlapping chunks: $\mathcal{C}(d_{\max}, 1), \dots, \mathcal{C}(d_{\max}, 2^{d_{\max}})$. Let $\mathcal{O}(k, \alpha_1, \alpha_2)$ represent the k^{th} overlapping chunk. $\mathcal{O}_k(\alpha_1, \alpha_2)$ is the union of the trailing sub-chunk of $\mathcal{C}(d_{\max}, k)$ at split factor α_1 and the leading sub-chunk of $\mathcal{C}(d_{\max}, k + 1)$ at split factor α_2 . We have $\mathcal{O}_k(\alpha_1, \alpha_2) = \mathcal{T}(\mathcal{C}(d, k), \alpha_1) \cup \mathcal{L}(\mathcal{C}(d, k + 1), \alpha_2)$ where $k \in \{1, 2, \dots, 2^d - 1\}$.

A special case is the proportional overlap chunk $\mathcal{O}_k(\alpha)$. This is equivalent to $\mathcal{O}_k(\alpha_1, \alpha_2)$ with $\alpha_1 = 1 - \alpha$ and $\alpha_2 = \alpha$. The trailing sub-chunk of $\mathcal{C}(d, k)$ and the leading sub-chunk of $\mathcal{C}(d, k + 1)$ have similar paragraph proportions relative to their respective chunks. We have $\mathcal{O}_k(\alpha) = \mathcal{O}_k(1 - \alpha, \alpha) = \mathcal{T}(\mathcal{C}(d, k), 1 - \alpha) \cup \mathcal{L}(\mathcal{C}(d, k + 1), \alpha)$.

Algorithm 2 Overlapping Chunking**Require:** Text $\mathcal{X}$ with N paragraphs, maximum depth $d_{\max}$, split factors $\alpha_1, \alpha_2 \in [0, 1]$ **Ensure:** Non-overlapping chunks $\{\mathcal{C}(d_{\max}, k) : k = 1, \dots, 2^{d_{\max}}\}$ and overlapping chunks $\{\mathcal{O}_k(\alpha_1, \alpha_2) : k = 1, \dots, 2^{d_{\max}} - 1\}$

```

1: SPLIT(1, N, 1, 1)                                ▷ Step 1: Recursive binary splitting

2: function SPLIT(lo, hi, d, k)
3:    $\mathcal{C}(d, k) \leftarrow \mathcal{X}[\text{lo}:\text{hi}]$ 
4:   if  $d = d_{\max}$  then
5:     return
6:   end if
7:    $\text{mid} \leftarrow \lfloor (\text{lo} + \text{hi})/2 \rfloor$ 
8:   SPLIT(lo, mid, d + 1, 2k - 1)
9:   SPLIT(mid + 1, hi, d + 1, 2k)
10: end function

11: for  $k = 1$  to  $2^{d_{\max}} - 1$  do                ▷ Step 2: Construct overlapping chunks
12:    $N(d_{\max}, k) = |\mathcal{C}(d_{\max}, k)|$ 
13:    $N(d_{\max}, k + 1) = |\mathcal{C}(d_{\max}, k + 1)|$ 
14:    $t \leftarrow \lfloor (1 - \alpha_1) \cdot N(d_{\max}, k) \rfloor$ 
15:    $\ell \leftarrow \lceil \alpha_2 \cdot N(d_{\max}, k + 1) \rceil$ 
16:    $\text{trail}_k \leftarrow$  last  $t$  paragraphs of  $\mathcal{C}(d_{\max}, k)$ 
17:    $\text{lead}_{k+1} \leftarrow$  first  $\ell$  paragraphs of  $\mathcal{C}(d_{\max}, k + 1)$ 
18:    $\mathcal{O}_k(\alpha_1, \alpha_2) \leftarrow \text{trail}_k \cup \text{lead}_{k+1}$ 
19: end for

```

3.2 Local FCM Extraction

This focus on extracting an FCM for each text chunk with an LLM agent. Each process involves three sequential steps: (1) concept identification, distillation, and normalization, (2) causal validation, and (3) fuzzy edge weighting. The system prompt instructs the LLM agent to run these sequential steps. The agent then returns the corresponding local FCM.

Local Node Definition: The step takes in a text chunk and returns a list of nodes for its corresponding FCM. The process follows three sequential steps: *identification*, *distillation*, and *normalization*. Identification step finds every entity or idea that logically increases or decreases within the context of the text chunk. Distillation step reduces nouns by removing modifiers and attributive adjectives from all identified entities. This step merges synonymous entities into one and resolves pronouns to their respective noun antecedents.

Causal Validation: This step takes the list of nodes from the local node definition and filters out dead nodes. Dead nodes are entities that do not have a causal link to or from another entity. The causal validation process follows two sequential steps: *causal verification* and *causal discarding*. The causal verification

step checks every node in the list against the text chunk. The causal discarding step removes any remaining node that lacks a causal link within the context of their respective text chunk.

Fuzzy Edge Weighting: This step takes in a list of causally validated nodes and their causal relationships and returns a weighted edge list for its corresponding FCM. The process follows two sequential steps: *direction assignment* and *intensity scoring*. The direction assignment step determines the polarity of each relationship. It assigns a positive weight when the cause and effect move in the same direction and a negative weight when they move in opposite directions. The intensity scoring step then assigns a numerical weight based on the linguistic strength of the relationship in the text chunk.

The LLM agent returns the local FCM $\mathcal{F}_d(k)$ for input chunk $\mathcal{C}(d, k)$. This local extraction runs over chunks $\mathcal{C}(d, 1), \dots, \mathcal{C}(d, 2^d)$. Their respective local FCMs are $\mathcal{F}_d(1), \dots, \mathcal{F}_d(2^d)$. We now mix these local FCMs into a global FCM.

3.3 Node Consolidation

This process combines the node lists of multiple local FCMs into a single global node list. The process follows three sequential steps: (1) global deduplication, (2) thematic clustering, and (3) intra-theme semantic merging.

Global Deduplication: The step combines the node lists of all local FCMs into a single global list. This list may contain duplicate nodes. The LLM agent then scans the full list and eliminates redundancy at two levels. Exact string matches merge directly into a single surviving node. Near-duplicate nodes that share the same underlying meaning despite minor lexical differences also merge.

Thematic Clustering: The step takes the deduplicated node list and organizes it into high-level themes. It analyzes the semantic domain of each node and groups related nodes under a common theme. Every node must belong to exactly one theme.

Intra-theme Semantic Merging: The step examines the nodes within each theme and identifies those that represent different facets of the same underlying causal concept. It groups these nodes into semantic clusters and merges each cluster into a single consolidated node. This reduces the total number of nodes in the global FCM and constitutes a form of “dimensionality reduction”. Every node maps to exactly one consolidated node and no node is discarded.

The system prompt for node consolidation instructs the LLM agent to run global deduplication, thematic clustering, and intra-theme semantic merging. The agent returns a unified list of nodes for a set of local FCMs

Bayesian De-Chunking: The LLM-extracted FCMs from the text chunks mix through convex combination to give the mixed FCM $\mathcal{F}_{\text{Likelihood}}$ as described

above. These likelihood FCMs also give their corresponding posterior FCMs $\mathcal{P}_k$ via equation (5). The posterior FCMs may not have the same nodes or edges as their corresponding likelihood FCMs.

4 Simulation Results

We applied our chunking method to “The Thucydides Trap” article by Graham Allison [1]. So $\mathcal{X}$ represents the article (full-text). This article has 29 paragraphs.

We use the Gemini 3.1 Pro Preview [4] as the LLM agent for this simulation. Gemini 3.1 Pro Preview excels at complex problem-solving tasks that require advanced reasoning. We set the temperature to 0, top p to 1, and top K to 1.0 for all the Gemini API calls. We compared the effect of $d_{\max}$ on FCM extractions with our non-overlapping chunking (Algorithm 1). We also compared the effect of $d_{\max}$ and splitting factor α for proportional chunking (Algorithm 2).

We observed various FCM parameters in our simulations and compared how they respond to text chunking. Merged nodes refers to the collection of local node lists from a given set of text chunks and may include repeated nodes. Semantic themes are the categories that the thematic clustering step produces. Consolidated nodes are the nodes that the intra-theme semantic merging step generates. Attractors are the stable long-term states or recurring cycles that an FCM converges to after repeated iterations of its update rule.

We also observed the number of nonzero eigenvalues of the FCM weight matrices. All zero eigenvalues imply a degenerate causal structure where the FCM captures no meaningful causal relationships and converges to the trivial null attractor. Nonzero eigenvalues indicate that the FCM captures genuine causal relationships and produces richer and more informative dynamics.

Text Chunking: We tested our non-overlapping chunking method (Algorithm 1) and compared it to a baseline that feeds the entire text directly into the LLM agent for FCM extraction. We considered $d_{\max} \in \{0, 1, 2, 3, 4\}$. The baseline corresponds to $d_{\max} = 0$. Table 1 compares the baseline FCMs with the ones that the non-overlapping chunking method produces.

LLM agent extraction more information from the Thucydides Trap article with the non-overlapping chunking. We noticed that the LLM agent captures more information as the number of chunks increases. The number of chunks equals $2^{d_{\max}}$ so it increases with $d_{\max}$. The number of merged nodes, unified nodes, and semantic themes increases with more chunks.

We also repeated the simulation for overlapping chunking. We tested the proportional overlapping chunks. We noticed the same trend. Table 2 shows the result for the overlapping text chunking.

De-chunking of mixed FCMs: Algorithm 2 split the 30-paragraph article into 3 chunks. The 1st chunk contained paragraphs 1-15 and produced the 10-node likelihood FCM $\mathcal{F}_1$. The 2nd chunk contained paragraphs 7-23 and produced the 9-node likelihood FCM $\mathcal{F}_2$. The 3rd chunk contained paragraphs 15-30 and

Table 1: Thucydides trap and non-overlapping text chunking: FCM extraction with an agentic LLM and non-overlapping text chunking

Count	$d_{\max} = 0$	$d_{\max} = 1$	$d_{\max} = 2$	$d_{\max} = 3$	$d_{\max} = 4$
Chunks	1	2	4	8	16
Merged nodes	10	18	28	52	75
Unified nodes	9	10	19	25	28
Semantic themes	4	4	5	5	5
Attractors	1	1	1	6	2
Nonzero edges	10	14	23	29	42
Nonzero eigenvalues	0	0	0	7	7

 Table 2: Thucydides Trap and overlapping text chunking: FCM extraction with an agentic LLM and the overlapping text chunking method (Algorithm 2). We used proportional overlapping chunks with split factor $\alpha = 0.5$.

Count	$d_{\max} = 0$	$d_{\max} = 1$	$d_{\max} = 2$	$d_{\max} = 3$
Chunks	1	3	7	15
Merged nodes	10	28	51	96
Unified nodes	9	11	19	32
Semantic themes	4	5	6	6
Attractors	1	1	2	2
Nonzero edges	10	16	32	49
Nonzero eigenvalues	0	0	3	10

produced the 6-node likelihood FCM $\mathcal{F}_3$. The 11-node mixed-FCM $\mathcal{F}_{\text{Likelihood}}$ mixed $\mathcal{F}_1$, $\mathcal{F}_2$, and $\mathcal{F}_3$ with equal convex weights $\frac{1}{3}$. $\mathcal{F}_1$, $\mathcal{F}_2$, and $\mathcal{F}_3$ also produced their respective posterior FCMs: $\mathcal{P}_1$ with 8 nodes, $\mathcal{P}_2$ with 8 nodes, and $\mathcal{P}_3$ with 7 nodes. The 3 posterior FCMs then also mixed with equal mixing weights to give the 9-node mixed posterior-FCM $\mathcal{F}_{\text{Posterior}}$.

We asked the same question to all 8 FCMs: “What if the rising power’s ambition remains unchecked?” We ask this question by clamping ‘on’ the node “Rising Power’s Ambition & Entitlement”. The FCMs F_1 and F_2 converged to fixed points while their posterior FCMs $\mathcal{P}_1$ and $\mathcal{P}_2$ converged to limit cycles. $\mathcal{P}_1$ converged to a 2-step limit cycle while $\mathcal{P}_2$ converged to a 3-step limit cycle. But the “Likelihood of War” node stayed on in all 4 equilibria. So all 4 FCMs agreed that there would be war.

The FCM $\mathcal{F}_3$ disagreed. It converged to a fixed point where the “Likelihood of War” node stayed off. But its posterior FCM $\mathcal{P}_3$ converged to another fixed point where the “Likelihood of War” node stayed on. The likelihood-FCM mixture $\mathcal{F}_{\text{Likelihood}}$ converged to the same fixed point as $\mathcal{F}_1$ while the posterior-FCM mixture $\mathcal{F}_{\text{Posterior}}$ converged to the same fixed point as $\mathcal{P}_3$. Both FCM-mixtures agreed that there would be war.

The posterior FCM $\mathcal{P}_3$ had 2 edge weights whose magnitude exceeded 1. The FCM did not change its equilibria even when those magnitudes were set to 1.

5 Conclusion

Controlled intelligent agents can map text to text chunks and then map the chunks to sparse causal edge matrices that define feedback fuzzy cognitive maps.




Fig. 6: Equilibrium answers to the clamping question: What if the rising power's ambitions remain unchecked? The time steps are along the x -axis and the nodes are along the y -axis. The purple nodes are inactive and the yellow nodes are active. The clamped-on nodes are in red. (a), (c), and (e) show the equilibria from the 3 likelihood FCMs $\mathcal{F}_1$, $\mathcal{F}_2$, and $\mathcal{F}_3$ and (g) shows the equilibrium from their equal weight mixture $\mathcal{F}_{\text{Likelihood}}$. (b), (d), and (f) show the equilibria from the 3 posterior FCMs $\mathcal{P}_1$, $\mathcal{P}_2$, and $\mathcal{P}_3$ and (h) shows the equilibrium from their mixture $\mathcal{F}_{\text{Posterior}}$. The last node "Likelihood of War" stays on in every equilibrium except (e).

The chunk matrices naturally mix or combine to produce a likelihood-like FCM that represents the feedback causal structure of the entire sampled document. The mixing weights must be convex and can depend on input values or otherwise vary with time. The mixing structure naturally gives an operator-level set of Bayesian posterior FCM matrices (which may require clipping or normalizing some edge values). These document FCMs can then mix to represent larger documents. They can also combine in hierarchies to store and fire documents in a structure large knowledge base.

References

1. Allison, G.: The thucydides trap: Are the u.s. and china headed for war? The Atlantic (Sep 2015), published September 24, 2015
2. Cooper, G.F.: The computational complexity of probabilistic inference using bayesian belief networks. *Artificial intelligence* **42**(2-3), 393–405 (1990)
3. Glykas, M.: Fuzzy cognitive maps: Advances in theory, methodologies, tools and applications, vol. 247. Springer (2010)
4. Google DeepMind: Gemini 3.1 pro. <https://deepmind.google/models/gemini/pro/> (March 2026), preview model. Google DeepMind
5. Kosko, B.: Differential Hebbian learning. In: AIP Conference proceedings. vol. 151, pp. 277–282. American Institute of Physics (1986)
6. Kosko, B.: Fuzzy cognitive maps. *International journal of man-machine studies* **24**(1), 65–75 (1986)
7. Kosko, B.: Hidden patterns in combined and adaptive knowledge networks. *International Journal of Approximate Reasoning* **2**(4), 377–393 (1988)
8. Osoba, O., Kosko, B.: Beyond dags: modeling causal feedback with fuzzy cognitive maps. arXiv preprint arXiv:1906.11247 (2019)
9. 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)
10. Panda, A.K., Adigun, O., Kosko, B.: The agentic leash: Extracting causal feedback fuzzy cognitive maps with llms. arXiv preprint arXiv:2601.00097 (2025)
11. Panda, A.K., Adigun, O., Kosko, B.: Causal autoencoder-like generation of feedback fuzzy cognitive maps with an llm agent. In: 24th IEEE International Conference on Machine Learning and Applications (ICMLA). pp. 1234–1241. IEEE (2025)
12. 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)
13. 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)
14. 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)
15. 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)