



REVIEW

Fuzzy Cognitive Maps: A Comprehensive Survey and Comparative Analysis

Maikel Leon^{*}

Department of Business Technology, Miami Herbert Business School, University of Miami, Coral Gables, FL, USA

*Corresponding Author: Maikel Leon. Email: mleon@miami.edu

Received: 25 June 2026; Accepted: 20 August 2026; Published: 22 October 2026

ABSTRACT: Fuzzy cognitive maps (FCMs) are a soft computing paradigm that combines fuzzy logic, neural networks, and graph theory to represent causality and reasoning in complex systems. Since their introduction in the 1980s as a generalization of cognitive maps, they have matured into a versatile framework that spans many variants, learning algorithms, and application areas. This survey consolidates that body of work into a single structured reference. The underlying literature search was conducted and is reported with reference to the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) 2020 statement, with a full flow diagram and a reproducible per-database search protocol. It traces the development of FCMs from graph-based encodings of expert knowledge to modern modeling and decision-support tools, and it presents the underlying mathematics, including concept activation dynamics, weight update rules, activation and membership functions, defuzzification, and error functions for supervised learning. The survey positions itself against recent reviews, distinguishes the main reasoning-rule families (classical, memory-based, quasi-nonlinear, and long-term cognitive networks), and organizes Hebbian, gradient-based, evolutionary, swarm intelligence, reinforcement, and hybrid learning algorithms into a unified taxonomy, with a qualitative and critical comparison of their supervision requirements, computational complexity, scalability, and the assumptions under which each has been validated. The survey catalogs variants developed to handle uncertainty and richer dynamics, organized into uncertainty-oriented, temporal, and machine-learning-integrated groups, including gray, interval, neutrosophic, rough, high-order, random, dynamic, deep, quantum-inspired, and chaotic FCMs. It reviews stability and convergence criteria and open theoretical questions. Further sections cover evaluation metrics and datasets; interpretability across static, dynamic, causal-effect, and actionable (counterfactual) levels, together with explanation evaluation, fairness considerations, and expert-bias mitigation; a comparison with Bayesian networks, neural networks, system dynamics, and agent-based models, together with an explicit account of failure cases and boundary conditions under which those alternatives are preferable; practical guidelines for concept selection and membership-function tuning; applications across engineering, energy, health care, economics, and other domains; and available software tools. The survey closes by identifying the challenges and research directions that connect FCMs with modern machine learning and explainable artificial intelligence (XAI), and it is intended to serve as both a rigorous entry point for newcomers and a working reference for practitioners.

KEYWORDS: Fuzzy cognitive maps; causal modeling; learning algorithms; uncertainty modeling; explainable artificial intelligence; decision support systems; stability and convergence; soft computing

1 Introduction

Fuzzy cognitive maps (FCMs) occupy an unusual position in computational modeling: they are simple enough to be sketched by hand from expert judgment, yet expressive enough to simulate the nonlinear, feedback-driven behavior of complex systems. This pairing of transparency with dynamic reasoning explains

why FCMs have spread across fields as different as engineering, medicine, ecology, economics, and the social sciences. That same reach, however, has produced a literature that is rich but poorly organized. Foundational results, dozens of learning algorithms, a steadily growing family of model variants, and applications scattered across many domains are reported in venues that seldom reference one another, and most existing reviews examine only a single slice: one class of algorithms, a particular variant, or a narrow application area. Readers who want to understand how the pieces fit together have had nowhere to turn. This survey is written to close that gap.


This article provides a structured, comprehensive survey of FCMs that connects theory, methods, and practice in one place. On the theoretical side, it assembles the mathematical foundations of FCMs, additional activation and membership-function formulations, a taxonomy of learning algorithms with a qualitative performance comparison, a catalog of model variants for handling uncertainty and dynamics, and stability and convergence results. On the practical side, it covers evaluation metrics and datasets, interpretability considerations, implementation guidelines, applications across numerous domains, and a review of software tools, before closing with open challenges and future directions. A dedicated comparison contrasts FCMs with Bayesian belief networks, neural networks, system dynamics, and agent-based models, clarifying when each representation is most appropriate.

1.1 Survey Methodology and Contributions

The literature search and study selection underlying this survey were conducted and are reported with reference to the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) 2020 statement; the protocol was not registered. The literature was searched between January and May 2026 across Web of Science, Scopus, IEEE Xplore, the ACM Digital Library, and Google Scholar, using the query template (“fuzzy cognitive map*”) AND (“learning” OR “algorithm*” OR “variant*” OR “stability” OR “convergence” OR “application*”), adapted to the syntax of each database and covering items published between 1986 (the year FCMs were first introduced) and May 2026. [Table 1](#) reports, for each database, the exact query executed, the fields searched, the filters applied, and the number of records retrieved, so that the search can be reproduced. The searches identified 1100 records in total (Web of Science: 342, Scopus: 315, IEEE Xplore: 198, ACM Digital Library: 84, Google Scholar: 161). Removing 220 duplicates by DOI and normalized-title matching left 880 records for title-and-abstract screening against three inclusion criteria (peer-reviewed or formally archived; proposing a novel theory, method, variant, or substantive application of FCMs; written in English) and two exclusion criteria (purely derivative restatements of earlier work; anecdotal reports without methodological content). Screening excluded 640 records, of which 132 were removed for not being written in English. The full texts of the remaining 240 reports were all retrieved and assessed for eligibility; 137 were excluded at this stage (71 as derivative restatements of earlier work, 45 for lacking methodological content, 18 as anecdotal reports without a substantive FCM contribution, 1 report that was initially included but removed during revision after being retracted by its publisher, and 2 reports found during a revision audit to fall outside the inclusion criteria), leaving 103 studies included through the database stream. Backward citation tracing of the included studies yielded 9 additional reports, including foundational works that predate the 1986 search window, and 4 further reports satisfying the same eligibility criteria were identified and incorporated while revising this article, bringing the total to 116 studies in the corpus synthesized in this survey. Two additional references, a mathematical monograph on the joint spectral radius cited as background in [Section 7](#) and a standard reinforcement-learning textbook, lie outside the FCM corpus, so the bibliography comprises 118 entries in total. [Fig. 1](#) presents the corresponding PRISMA 2020 flow diagram. Following the classification scheme of Schuerkamp and Giabbanelli’s systematic review [1], the collected works were grouped into categories covering mathematical foundations, learning algorithms, stability and convergence analysis, model variants, applications, and software tools.

Table 1: Reproducible search protocol: query executed in each database, fields searched and filters applied, and records retrieved. Search window: 1986–May 2026; searches executed January–May 2026.

Database	Query as Executed	Fields and Filters	Records
Web of Science (Core Collection)	TS = (“fuzzy cognitive map*” AND (“learning” OR “algorithm*” OR “variant*” OR “stability” OR “convergence” OR “application*”))	Topic; 1986–2026; articles, reviews, proceedings	342
Scopus	TITLE-ABS-KEY (“fuzzy cognitive map*” AND (learning OR algorithm* OR variant* OR stability OR convergence OR application*))	Title, abstract, keywords; 1986–2026	315
IEEE Xplore	(“fuzzy cognitive map” AND (learning OR algorithm OR variant OR stability OR convergence OR application))	Metadata; 1986–2026	198
ACM Digital Library	[[Title: “fuzzy cognitive map*”]] AND [[Abstract: learning OR algorithm* OR variant* OR stability OR convergence OR application*]]	Title and abstract; 1986–2026	84
Google Scholar	“fuzzy cognitive map” learning OR algorithm OR variant OR stability OR convergence OR application	Relevance-ranked export; 1986–2026	161

**Figure 1:** PRISMA 2020 flow diagram of the study identification, screening, and inclusion process underlying this survey. Records were identified from five databases, through backward citation searching of included studies, and during the revision of this article; one initially included report was removed after its retraction, and two were removed as falling outside the inclusion criteria. The final corpus comprises 116 studies. The bibliography also contains two background references outside this corpus (a mathematical monograph on the joint spectral radius and a standard reinforcement-learning textbook), so the reference list totals 116 + 2 = 118 entries.

This survey makes several distinct contributions beyond existing reviews. First, it provides a unified taxonomy of FCM learning algorithms and model variants, highlighting how each addresses specific limitations of the classical FCM. Second, it discusses concept selection and data integration strategies for linking numerical variables to abstract FCM concepts, and it offers practical guidelines for choosing and tuning fuzzy membership functions. Third, it compares learning algorithms qualitatively in terms of supervision requirements, computational complexity, and scalability, drawing on trends reported in the literature. Fourth, it synthesizes recent theoretical advances in stability analysis and identifies open research questions. Fifth, it critically analyzes emerging variants such as deep, quantum-inspired, and chaotic FCMs, clarifying their fundamental differences, advantages, and challenges, and evaluating them explicitly against the interpretability framework developed later in the survey. Sixth, it delineates failure cases and boundary conditions under which FCMs are demonstrably outperformed by alternative frameworks, and it distills a cross-cutting synthesis that takes positions on the main tensions among competing reasoning rules, learning algorithms, and variants. Finally, it outlines future research directions that integrate FCMs with modern machine learning and explainable artificial intelligence (XAI), thereby situating FCMs within the evolving landscape of computational intelligence. Combined with a broad account of applications and tools, the survey aims to be an authoritative, forward-looking reference for researchers and practitioners alike.

The scope of this survey has limits worth stating plainly. The search was confined to English-language sources indexed in the consulted databases, so relevant work in other languages or in venues outside those indexes may be underrepresented. A single author performed screening and eligibility assessment, so inter-rater reliability statistics cannot be reported; this is mitigated by the documented criteria and the reproducible protocol of [Table 1](#). The synthesis is narrative, and no meta-analysis or pooled effect size is reported. Quantitative results cited from individual studies are reproduced as reported by their authors and were not independently re-evaluated, since the absence of shared FCM benchmarks prevents a controlled head-to-head comparison.

1.2 Positioning with Respect to Recent Surveys

Several valuable reviews of the FCM literature have appeared recently, and it is fair to ask what the present survey adds. Schuerkamp and Giabbanelli's systematic review [\[1\]](#) catalogs 26 FCM extensions proposed between 2011 and March 2023 under a rigorous protocol, yet its scope is deliberately confined to extensions rather than the full modeling pipeline. Nápoles et al. [\[2\]](#) review recent advances and algorithms, emphasizing pattern classification and time-series forecasting and treating FCMs as recurrent neural systems; their algorithmic coverage is deep, while applications, software, and human factors receive comparatively less attention. Apostolopoulos and Groumpos [\[3\]](#) examine specifically the role of FCMs within XAI. Most recently, Sovatzidi and Iakovidis [\[4\]](#) reviewed FCMs and their perspectives for explainable decision making, consolidating the case for FCMs as inherently interpretable decision-support models and mapping the routes through which they can deliver explanations; their review is the closest in spirit to the interpretability treatment of [Section 9](#), and the levels-of-explanation reading adopted there is consistent with the perspectives they articulate. The present survey complements these works: it spans the pipeline from mathematical foundations, reasoning rules, and learning algorithms through variants, stability theory, evaluation, interpretability, comparisons with other frameworks, implementation guidance, applications, and software in a single self-contained treatment. [Table 2](#) summarizes the comparison.

Table 2: Coverage comparison between the present survey and recent FCM reviews. “Full” denotes in-depth coverage, “Partial” denotes limited coverage, and a dash denotes minimal or no coverage.

Aspect	Schuerkamp & Giabbanelli [1]	Nápoles et al. [2]	Apostolopoulos & Groumpos [3]	Sovatzidi & Iakovidis [4]	This Survey
Foundations and reasoning rules	Partial	Full	Partial	Partial	Full
Learning algorithms	Partial	Full	–	Partial	Full
Variants and extensions	Full	Partial	–	Partial	Full
Stability and convergence theory	–	Partial	–	–	Full
Evaluation metrics and datasets	–	Partial	–	–	Full
Interpretability and human factors	–	Partial	Full	Full	Full
Comparison with other frameworks	–	–	Partial	–	Full
Implementation guidance and software	–	Partial	–	–	Full
Application domains	Partial	–	Partial	Partial	Full

1.3 Reading Map

Table 3 provides a reading map: each entry summarizes the purpose of a section and suggests a natural order in which to explore the topics. Readers primarily interested in theory may focus on the mathematical formulation, additional formulations, learning algorithms, and stability sections, whereas practitioners may proceed directly to the implementation guidelines and application domains. The comparison with other modeling frameworks serves readers who are still deciding whether an FCM is the right instrument for their problem.

Table 3: Reading map summarizing the structure of the survey and suggesting an order for readers to follow according to their interests.

Section	Purpose
Origins and historical development	Traces the roots of cognitive mapping and fuzzy logic and the emergence of the first FCMs
Mathematical formulation	Defines concepts and weights, the activation dynamics and reasoning-rule families that update them, and the error functions used in supervised learning
Additional formulations	Presents further activation and fuzzy membership functions with defuzzification, and gray, rough, neutrosophic, high-order, random, dynamic, and reinforcement-learning formulations

(Continued)

Table 3 (continued)

Section	Purpose
Learning algorithms	Organizes Hebbian, evolutionary, swarm intelligence, reinforcement, and hybrid methods into a taxonomy and compares their empirical performance
Variants and extensions	Surveys gray, interval, neutrosophic, rough, high-order, random, dynamic, deep, quantum-inspired, and chaotic FCMs
Stability and convergence	Reviews convergence and stability criteria, including spectral and Lyapunov conditions, and the questions that remain open
Evaluation metrics and datasets	Introduces performance metrics for classification and regression and lists representative datasets
Interpretability and human factors	Examines interpretability at the static, dynamic, causal-effect, and actionable (counterfactual) levels, together with explanation evaluation, fairness, similarity measures, and expert-bias mitigation
Comparison with other modeling frameworks	Contrasts FCMs with Bayesian networks, neural networks, system dynamics, and agent-based models, and delineates failure cases and boundary conditions
Implementation guidelines and best practices	Gives practical steps for concept selection, data integration, weight assignment, and model evaluation
Application domains	Reviews applications in travel behavior, pandemic modeling, engineering, energy, sustainability, health care, economics, and computing, with illustrative case studies
Software tools	Reviews available software and the features each provides for building and analyzing FCMs
Cross-cutting synthesis	Takes positions on the recurring tensions: convergence vs. discrimination, learner choice, interpretability claimed vs. delivered, and variant proliferation vs. consolidation
Challenges and future research directions	Identifies open problems in scalability, uncertainty handling, stability, integration with modern artificial intelligence (AI), and ethical considerations

The remainder of this article is organized as follows. [Section 2](#) traces the origins and historical development of FCMs. [Section 3](#) formalizes their mathematical foundations, and [Section 4](#) collects additional activation, membership, and uncertainty-oriented formulations. [Section 5](#) organizes the learning algorithms into a taxonomy and compares their empirical behavior, while [Section 6](#) surveys the model variants and extensions. [Section 7](#) reviews stability and convergence results, [Section 8](#) covers evaluation metrics and datasets, and [Section 9](#) examines interpretability and human factors. [Section 10](#) compares FCMs with alternative modeling frameworks and delineates failure cases; [Section 11](#) offers implementation guidelines and best practices; [Section 12](#) reviews application domains with in-depth case studies; and [Section 13](#) surveys the software ecosystem. [Section 14](#) distills a cross-cutting synthesis of the assembled evidence, [Section 15](#) discusses open challenges and future research directions, and [Section 16](#) concludes the article.

2 Origins and Historical Development

Fuzzy cognitive maps did not emerge in isolation; they grew out of two older intellectual traditions that matured side by side. This section traces those roots, beginning with research in psychology and political science that used directed graphs to represent causal beliefs, moving through the integration of fuzzy set theory to handle vagueness, and culminating in the first formulations of FCMs in the 1980s. The section also introduces the contributors who later extended the theory and whose work on modeling, learning, convergence analysis, and related topics broadened the methodological and application landscape.

2.1 Cognitive Maps and Fuzzy Logic

The concept of a cognitive map originates in psychology and political science, where researchers sought to model how individuals perceive causal relationships among the factors of a complex environment. Axelrod's influential work used signed directed graphs to represent the beliefs of political decision-makers, emphasizing that the structure of such a map shapes the conclusions its holder can reach [5]. In a separate line of inquiry, Zadeh had introduced fuzzy set theory to describe concepts whose boundaries are not sharply defined [6]. Fuzzy sets assign each element a membership degree in the interval $[0, 1]$, enabling the representation of linguistic variables such as “high”, “medium”, and “low”. Together, these developments laid the groundwork for combining cognitive maps with fuzzy logic.

In the mid-1980s, Kosko generalized Axelrod's cognitive maps by allowing causal relationships to take fuzzy values. In his foundational paper on FCMs [7], he proposed representing causal influences by weights in the continuous interval $[-1, 1]$, where positive values denote promoting influences and negative values denote inhibiting influences. This innovation permitted the modeling of ambiguous and graded causation, bridging the gap between qualitative reasoning and quantitative modeling.

2.2 Emergence of FCMs

The first decade of FCM research focused on establishing mathematical foundations and exploring potential applications. Kosko's early papers investigated properties of FCM dynamics and introduced the adaptive inference paradigm [8]. Styblinski and Meyer introduced FCMs into qualitative circuit analysis, comparing them with signal flow graphs and examining how the maps converge to equilibria [9]. Taber demonstrated the usefulness of FCMs for representing knowledge and processing uncertain information [10], and Dickerson and Kosko used them to animate virtual worlds, showcasing their ability to capture feedback loops in complex systems [11]. By the mid-1990s, FCMs were established as a natural tool for modeling causal relations in social, economic, and technical systems, a tradition that Kosko and collaborators later carried into policy analysis with maps of public support for insurgency and terrorism [12].

Fig. 2 illustrates the structure shared by these early models, namely a small set of concepts linked by signed, weighted arrows that can form feedback loops.

Review articles in subsequent decades have charted the growth of FCM research, and comprehensive treatments of the basic theory and applications appeared in the early 2010s. Groumpos's mathematical description of FCM models introduced new construction methods and algorithms, demonstrating their applicability through examples from manufacturing and health decision-making [13]. Papageorgiou and Salmeron charted a full decade of FCM research [14]; Papageorgiou reviewed the learning algorithms proposed up to that point [15]. A recent study surveyed engineering applications, insights, and future directions [16]. A comprehensive review of FCMs in systems risk analysis cataloged further modeling advances while highlighting persistent challenges of scalability and interpretability [17], and a systematic review in *ACM Computing Surveys* cataloged 26 extensions of FCMs, observing that few address multiple

limitations at once and that publicly available code is scarce [1]. Recent work increasingly connects FCMs with machine learning and deep learning techniques [18–20]. The historical trajectory of FCMs thus runs from manually drawn cognitive maps to sophisticated computational models that learn from data and adapt to dynamic environments.




Figure 2: A simple fuzzy cognitive map relating lifestyle factors to physical fitness. Each node is a concept, and each directed arrow encodes a causal influence whose sign and magnitude give the direction and strength of that influence. Solid arrows denote positive (promoting) influence and dashed arrows denote negative (inhibiting) influence, with weights in the interval $[-1, 1]$. The two arrows between sleep quality and stress form a feedback loop, a characteristic feature of FCMs.

3 Mathematical Formulation

This section formalizes the key mathematical constructs underlying FCMs: representing concepts and weights as vectors and matrices, the iterative rules that update concept activations, and the error functions used in supervised learning. These definitions carry the weight of everything that follows, since the learning algorithms, model variants, and stability analyses of later sections are all stated in the language introduced here.

3.1 Structure and Representation

An FCM consists of a set of n concepts $C = \{C_1, C_2, \dots, C_n\}$, represented as the nodes of a directed graph, together with weighted directed edges encoding causal influences. A weight $w_{ij} \in [-1, 1]$ is assigned to each ordered pair (C_j, C_i) , indicating the degree to which concept C_j influences concept C_i . Positive weights signify that an increase in C_j tends to increase C_i , whereas negative weights signify inhibitory effects. The treatment of self-loops merits clarification. In the original formulation of FCMs, the diagonal of the weight matrix was often set to zero, effectively assuming that concepts do not influence themselves. However, subsequent work has shown that nonzero diagonal entries can capture inertia or memory effects in the dynamics, and participatory mapping studies explicitly accommodate self-loop weights when interviewees describe feedback within a single variable [21,22]. Other perspectives explicitly exclude self-loops because FCMs are weighted directed networks without diagonal edges [13]. Whether to include self-loops therefore depends on the modeling goals: setting $w_{ii} = 0$ yields a purely interconceptual map, whereas assigning $w_{ii} \neq 0$ introduces self-reinforcement or decay. The weights are arranged into an adjacency matrix $W = [w_{ij}]$, and the state of the FCM at discrete time k is described by an activation vector $\mathbf{A}(k) = (A_1(k), \dots, A_n(k))$. Table 4 collects the principal symbols used throughout the survey.

Table 4: Principal notation used throughout the survey.

Symbol	Meaning
n	Number of concepts in the map
C_i	The i -th concept (node)
w_{ij}	Causal weight of concept C_j on C_i , $w_{ij} \in [-1, 1]$
$W = [w_{ij}]$	Weight (adjacency) matrix
$\mathbf{A}(k)$	Activation vector at iteration k
$\mathbf{A}^*$	Fixed-point (equilibrium) activation vector
f	Bounded, non-decreasing activation function
λ	Gain (steepness) parameter of f
η	Learning rate
$\rho(W)$	Spectral radius of W
$\ W\ $	Induced (operator) norm of W
L	Lipschitz constant of the activation function f
L_{CE}	Cross-entropy loss (Section 11)
E	Training error
m	Number of training patterns, indexed by p

3.2 Concept Activation Dynamics

Iterative update rules govern the evolution of concept activations. The most common update scheme computes a weighted sum of incoming influences and applies a squashing function to keep activations within a bounded interval. In vector form,

$$\mathbf{A}(k+1) = f(W \mathbf{A}(k)), \quad (1)$$

where f is a bounded, nondecreasing activation function applied elementwise. Common choices include the hyperbolic tangent, logistic, and linear saturation functions. For example, the sigmoid function

$$f(x) = \frac{1}{1 + \exp(-\lambda x)} \quad (2)$$

maps real values to $(0, 1)$, with gain parameter $\lambda > 0$, whereas the hyperbolic tangent function

$$f(x) = \tanh(\lambda x) \quad (3)$$

maps values to $(-1, 1)$ and accommodates symmetric positive and negative activations. The choice of f influences both the stability and the interpretability of the FCM [[9,23](#)].

The update can also be written componentwise,

$$A_i(k+1) = f\left(A_i(k) + \sum_{j=1}^n w_{ij} A_j(k)\right), \quad (4)$$

where retaining $A_i(k)$ in the sum introduces a memory (inertia) term equivalent to a unit self-loop, and dropping it recovers the memoryless update of [Eq. \(1\)](#). The iteration continues until the activation vector reaches a fixed point $\mathbf{A}^*$ satisfying $\mathbf{A}^* = f(W \mathbf{A}^*)$, a limit cycle, or a chaotic attractor [[7,24](#)]. Which of these behaviors emerges depends jointly on the weights and the activation function: for the linear update

without squashing, a spectral radius $\rho(W) < 1$ drives the state to the origin, while for nonlinear activations, convergence to a unique fixed point is usually established through the contraction conditions reviewed in Section 7. Nonlinear activations combined with large weight magnitudes can instead produce oscillations or chaotic dynamics.

3.3 Reasoning-Rule Families

Eqs. (1) and (4) define what is often called the classical reasoning rule, yet several families of reasoning rules now coexist in the literature, and the choice among them affects convergence, expressiveness, and interpretability. Four families deserve explicit mention. First, the classical rule of Kosko [7] computes each activation as the weighted sum of the other concepts, optionally followed by a squashing function. Second, modified rules with memory retain the previous activation $A_i(k)$ in the sum, equivalent to a unit self-loop, and introduce inertia; rescaled versions average the memory and network terms to keep activations bounded. Third, quasi-nonlinear reasoning, introduced with quasi-nonlinear FCMs, interpolates between the nonlinear update and the initial stimulus via a nonlinearity coefficient, which provably prevents the map from collapsing to a unique fixed point that ignores the input [25]. Fourth, long-term cognitive networks and their recurrence-aware successors relax the requirement that weights be confined to $[-1, 1]$ and learn inner concept transformations, trading part of the causal reading of individual weights for markedly better predictive performance in classification and forecasting tasks [26,27]. A practical reading of this progression is that each family relaxes one constraint of the classical rule (memorylessness, guaranteed uniqueness of attractors, or bounded causal weights) in exchange for expressiveness. The modeler should adopt the weakest relaxation the task requires, since each step away from the classical rule also weakens the direct causal interpretation of the weight matrix.

3.4 Error Functions and Supervised Learning

When training an FCM from data, one can define a loss function that measures the discrepancy between predicted and target activations. Suppose m training patterns are given, each with desired concept values $\mathbf{T}^{(p)} = (T_1^{(p)}, \dots, T_n^{(p)})$, where the superscript p indexes patterns and k remains reserved for iteration steps. The error function can be defined as

$$E = \frac{1}{2} \sum_{p=1}^m \sum_{i=1}^n (A_i^{(p)} - T_i^{(p)})^2, \quad (5)$$

where $A_i^{(p)}$ denotes the FCM output for pattern p after a sufficient number of iterations. Gradient descent updates the weights in the direction opposite to the gradient,

$$\Delta w_{ij} = -\eta \frac{\partial E}{\partial w_{ij}}, \quad (6)$$

with learning rate $\eta > 0$. While conceptually simple, gradient methods can be computationally expensive because the derivative must be unfolded through iterative updates over multiple time steps, as in backpropagation through time for recurrent networks. Alternative approaches, such as Hebbian learning and evolutionary algorithms, circumvent derivative computation altogether, as described in the following sections.

4 Additional Formulations

This section gathers additional equations and theoretical formulations that recur throughout the FCM literature, so that later sections can refer to them without interruption. Readers who prefer to move directly to the learning algorithms may treat this section as a reference and return to it as needed.

4.1 Additional Activation Functions

Besides the sigmoid and hyperbolic tangent functions of Eqs. (2) and (3), several other activation functions recur in FCM research; since they are standard piecewise-linear constructions, prose descriptions suffice. The linear-saturation function outputs zero for negative inputs, grows linearly across the unit interval, and saturates at one thereafter, confining activations to $[0, 1]$ without the sigmoid's smooth compression. The hard-limiting (threshold) function returns 1 for positive inputs and 0 otherwise, yielding binary activations appropriate for on/off concepts. The parametric rectified linear unit passes positive inputs through unchanged and scales negative inputs by a slope parameter $a \in (0, 1)$; unlike the preceding functions, it is unbounded, so it is typically paired with one of the weight normalization strategies described in Section 4.5. Choosing among these functions balances expressiveness against stability: bounded smooth functions favor convergence, piecewise-linear ones keep the mapping transparent, and unbounded ones demand explicit normalization.

4.2 Fuzzy Membership Functions and Defuzzification

In many FCM applications, the weights are obtained by translating linguistic expert judgments into numerical values using membership functions. A triangular membership function with parameters (a, b, c) is defined as

$$\mu_{\text{tri}}(x; a, b, c) = \begin{cases} 0, & x \leq a, \\ \frac{x-a}{b-a}, & a < x < b, \\ \frac{c-x}{c-b}, & b \leq x < c, \\ 0, & x \geq c, \end{cases} \quad (7)$$

while a trapezoidal membership function with parameters (a, b, c, d) is given by

$$\mu_{\text{trap}}(x; a, b, c, d) = \begin{cases} 0, & x \leq a, \\ \frac{x-a}{b-a}, & a < x < b, \\ 1, & b \leq x \leq c, \\ \frac{d-x}{d-c}, & c < x < d, \\ 0, & x \geq d. \end{cases} \quad (8)$$

The Gaussian membership function is defined as

$$\mu_{\text{gauss}}(x; c, \sigma) = \exp\left(-\frac{(x-c)^2}{2\sigma^2}\right), \quad (9)$$

where c and σ denote the center and the width (standard deviation). After obtaining membership degrees μ_u for candidate numerical values x_u , defuzzification can be performed via the weighted average

$$y = \frac{\sum_{u=1}^q x_u \mu_u}{\sum_{u=1}^q \mu_u}, \quad (10)$$

which yields a crisp numerical weight; this is the discrete form of centroid (center-of-gravity) defuzzification. The triangular, trapezoidal, and Gaussian shapes are the three most commonly encountered in practice, and they recur whenever expert knowledge is elicited for FCM construction.

Choosing and tuning membership functions. The choice of membership function shapes both the interpretability and the stability of the resulting FCM. Simple shapes such as triangular and trapezoidal functions are often preferred when expert assessments are expressed linguistically (e.g., “low”, “medium”, “high”) because the parameters (a, b, c, d) correspond directly to breakpoints on the linguistic scale. Gaussian functions provide smoother transitions and suit concepts whose activation varies continuously. Open-source tools such as the FCMpy library automatically generate and optimize membership functions over the universe of discourse using triangular, trapezoidal, and Gaussian shapes [28]. Parameters can be tuned via cross-validation to minimize prediction error or to maximize classification accuracy. When no prior knowledge is available, a common strategy is to discretize the universe of discourse into equidistant points and search for the membership parameters that minimize a loss function, subject to smoothness constraints. The selection should also anticipate the downstream defuzzification method: centroid-based defuzzification, as in Eq. (10), pairs well with symmetric membership functions, whereas ranking-based methods can accommodate asymmetric shapes. Table 5 distills these considerations into practical guidelines for common settings.

Table 5: Guidelines for choosing and optimizing fuzzy membership functions. These recommendations balance interpretability, computational cost, and data fidelity.

Membership Type	When to Use	Optimization Tips
Triangular	Linguistic assessments with clear lower, middle and upper points; low computational cost	Set (a, b, c) based on breakpoints; adjust b to fit median expert rating; optimize using cross-validation
Trapezoidal	Concepts exhibiting plateau regions (e.g., “normal” range)	Choose (a, b, c, d) so that $[b, c]$ corresponds to the plateau; tune edge slopes to control sensitivity
Gaussian	Smooth variations or noisy data	Initialize (c, σ) at the empirical mean and standard deviation; optimize via gradient-based search or evolutionary algorithms
Hybrid / adaptive	Unknown or complex distributions	Combine multiple shapes or use fuzzy neural networks that learn membership functions from data; regularize to prevent overfitting

4.3 Gray, Rough, and Neutrosophic Formulations

Gray FCMs represent each weight as an interval $[w_{ij}^{\min}, w_{ij}^{\max}]$. Interval arithmetic propagates through the update rule:

$$\tilde{A}(k+1) = f(\tilde{W} \otimes \tilde{A}(k)), \quad (11)$$

where $\otimes$ denotes interval matrix multiplication. Rough cognitive maps represent each concept as a pair of lower and upper approximations $\underline{C}_i$ and $\overline{C}_i$ determined by an indiscernibility relation. For a target set X , the lower approximation is defined as

$$\underline{X} = \{x \in U \mid B(x) \subseteq X\}, \quad (12)$$

and the upper approximation as

$$\overline{X} = \{x \in U \mid B(x) \cap X \neq \emptyset\}, \quad (13)$$

where $B(x)$ denotes the set of elements indiscernible from x and U is the universe of discourse. Neutrosophic cognitive maps assign to each weight a triplet $N = (T, I, F)$ of truth, indeterminacy, and falsity membership degrees with $T, I, F \in [0, 1]$ and no restriction on their sum. A neutrosophic weight update may be represented as

$$N'_{ij} = (T'_{ij}, I'_{ij}, F'_{ij}) = (f_T(T_{ij}, I_{ij}, F_{ij}), f_I(T_{ij}, I_{ij}, F_{ij}), f_F(T_{ij}, I_{ij}, F_{ij})), \quad (14)$$

where f_T, f_I, f_F are appropriate update functions. These formulations allow FCMs to model uncertainty and indeterminacy in expert knowledge. Theoretical analysis of convergence and stability for such extended maps remains an active research area [29–31].

4.4 High-Order, Random, and Dynamic Formulations

High-order FCMs generalize the standard model by including high-order interactions among concept activations. A second-order FCM update can be written as

$$A_i(k+1) = f\left(A_i(k) + \sum_{j=1}^n w_{ij} A_j(k) + \sum_{j=1}^n \sum_{l=1}^n w_{ijl} A_j(k) A_l(k)\right), \quad (15)$$

where w_{ijl} denotes the influence of the joint activation of C_j and C_l on C_i (with $j, l \in \{1, \dots, n\}$ and $w_{ijl} = w_{ilj}$ by symmetry of the pair) [32]. Random FCMs treat weights as random variables drawn from distributions such as

$$w_{ij} \sim \mathcal{N}(\mu_{ij}, \sigma_{ij}^2), \quad (16)$$

and compute expectations of concept activations via Monte Carlo simulation [33]. Dynamic FCMs introduce time-dependent weights $w_{ij}(t)$ updated according to differential or difference equations. A simple relaxation model,

$$w_{ij}(t+1) = \alpha w_{ij}(t) + (1 - \alpha) w_{ij}^0, \quad (17)$$

with $0 < \alpha < 1$, pulls each weight exponentially toward a baseline value w_{ij}^0 . The weight change may also respond to feedback from system outputs or reward signals [34]. Such dynamic updates enable FCMs to adapt to nonstationary environments.

4.5 Reinforcement Learning and Optimization Schemes

In reinforcement learning, weights are adjusted to maximize expected reward. The standard vehicle is temporal-difference learning of the Q-learning type: the estimated value of taking an action in a given state is nudged toward the sum of the immediate reward and the discounted value of the best successor action, with a learning rate governing the step size and a discount factor $\gamma \in [0, 1)$ weighting future rewards. The weight matrix is parameterized as a function of the learned action values, and gradient ascent adjusts the parameters accordingly. In genetic algorithms, candidate FCMs are ranked by a fitness function that, in the most common formulation, is simply the reciprocal of the training error of Eq. (5), with a small constant added to guard against division by zero; mutation and crossover operators then explore the search space. For particle swarm optimization (PSO), the velocity and position update equations are given in Section 5.4. After any of these updates, weight normalization keeps the weights within $[-1, 1]$; the schemes in use are global scaling (dividing every weight by the largest absolute weight in the matrix), row scaling (dividing each weight by the absolute sum of its row), smooth squashing through the hyperbolic tangent, and hard clipping at the interval bounds. As for hyperparameters, the ranges reported across the literature are conventional: learning rates between 10^{-4} and 10^{-1} , momentum up to 0.9, populations of 20 to 200 individuals for genetic and swarm methods, PSO inertia between 0.4 and 0.9 with cognitive and social coefficients between 1.0 and 2.5, and discount factors between 0.8 and 0.99; these values offer sensible starting points for new implementations.

5 Learning Algorithms

Learning algorithms adjust the weight matrix W so that the FCM reproduces observed system behavior or refines elicited expert knowledge. A tension runs through every method surveyed here: the more aggressively an algorithm reshapes the weight matrix to fit data, the harder it becomes to defend each weight as a causal judgment an expert would endorse. Keeping that tension in view helps explain why so many families of algorithms coexist.

5.1 Taxonomy of Learning Algorithms

Learning algorithms for FCMs can be broadly categorized by their underlying optimization principle and degree of supervision. Unsupervised Hebbian methods adjust weights based on co-activation of concepts [35,36]. In contrast, supervised and hybrid schemes employ evolutionary search such as genetic algorithms (GA) [37–39], swarm intelligence such as PSO [40,41], reinforcement learning (RL) signals [42,43], or gradient descent on a loss function [44,45]. This taxonomy helps researchers select suitable algorithms for a specific application, highlighting trade-offs among interpretability, computational cost, and scalability. Fig. 3 depicts the taxonomy, and Table 6 summarizes the main classes of learning algorithms used to train FCM weights, including hybrid neuro-fuzzy, convex-optimization, self-organizing, and hybrid evolutionary schemes [46–50].




Figure 3: Taxonomy of FCM learning algorithms by degree of supervision and underlying optimization principle. The four families correspond to the rows of Table 6 and are compared qualitatively at the end of this section.

Table 6: Major learning algorithms for FCMs.

Algorithm	Principle	Key References
Hebbian learning	Co-activation strengthens weights	[35]
Nonlinear Hebbian	Learning based on activation changes	[36]
Genetic algorithm	Evolutionary optimization of weights	[37–39]
Particle swarm optimization	Swarm search guided by local and global best	[39–41]
Memetic PSO	PSO with local search refinements	[41]
Reinforcement learning	Reward-driven adaptive updates	[42,43]
Gradient descent	Error minimization through unrolled iterations	[44,45]
Hybrid neuro-fuzzy	Neural networks integrated with FCM	[20,46]
Convex optimization	Direct fitting of high-order maps to data	[47]
Self-organizing	Node growth and pruning informed by self-organizing networks	[48]
Hybrid evolutionary	Combination of Hebbian and evolutionary	[49,50]

5.2 Empirical Performance across Learning Algorithms

Although many learning algorithms have been proposed for FCMs, their empirical performance varies across problem domains, dataset sizes, and noise levels. Early comparative studies found that GA and PSO outperformed simple Hebbian updates on small and medium-sized datasets, achieving lower error and faster convergence [39,40]. However, both require careful tuning of population size, inertia, and mutation parameters and may struggle with very large concept sets. Reinforcement learning approaches can adapt to nonstationary environments but incur substantial computational cost [42]. Deep FCMs, which embed FCM inference within neural networks, have reported accuracy improvements over the conventional baselines used in those specific studies (figures on the order of tens of percent in the most favorable cases), at the cost of reduced interpretability and longer training; such numbers are experiment-specific and should not be read as general trends [19,20]. Quantum-inspired evolutionary algorithms accelerate convergence and improve classification accuracy on benchmark datasets by exploiting superposition and interference effects [51]. At the same time, chaotic FCMs can capture irregular dynamics but are sensitive to noise and demand fine-grained parameter tuning [24]. A critical reading of this literature suggests the following: Hebbian methods work when the initial expert map is already close to correct and fail when substantial structural correction is needed; population-based methods perform well on small and medium maps but scale poorly and are rarely tested beyond a few dozen concepts; reinforcement approaches are compelling in nonstationary settings, but the evidence base is thin; and deep hybrids buy accuracy at a direct cost to the causal readability that motivates FCMs in the first place. The assumptions under which each family has been validated are narrow, and head-to-head evidence is weak because the studies use different datasets, metrics, and protocols. Hence, the comparison remains qualitative rather than a controlled benchmark. Table 7 summarizes the reported trends for representative learning algorithms; its entries should be read as directional findings synthesized from the cited works rather than as measurements from any single experiment.

Table 7: Qualitative trends in empirical performance of representative FCM learning algorithms. The values summarize findings from multiple studies and highlight trade-offs among accuracy, convergence speed, and interpretability.

Algorithm	Scenario	Observed Trend	References
Genetic algorithm	Small/medium datasets; classification	Lower error than Hebbian; sensitive to parameter tuning	[39]
Particle swarm	Small/medium datasets; regression	Faster convergence than GA; moderate robustness to noise	[40]
Reinforcement learning	Dynamic control problems	Adaptive but computationally intensive	[42]
Deep FCM	Image and time-series classification	Study-specific accuracy gains reported; reduced interpretability	[19,20]
Quantum-inspired evolutionary	Binary classification on UCI datasets	Improved accuracy; faster convergence	[51]
Chaotic FCM	Nonlinear dynamic systems	Captures irregular behavior; highly sensitive to parameter perturbation	[24]

5.3 Hebbian and Nonlinear Hebbian Learning

Hebbian learning is inspired by the neurobiological principle that active units strengthen their connections simultaneously. In the context of FCMs, weights are updated in proportion to the product of concept activations. The basic Hebbian update rule for weight w_{ij} is

$$\Delta w_{ij} = \eta A_i(k) A_j(k), \quad (18)$$

where $A_i(k)$ and $A_j(k)$ are the activations of concepts C_i and C_j at iteration k , and η is a small learning rate [35]. This rule strengthens the connection between concepts that are simultaneously active. Nonlinear and differential Hebbian variants incorporate activation changes or nonlinearity to enhance learning stability. One such rule updates weights according to

$$\Delta w_{ij} = \eta (A_i(k+1) - A_i(k)) A_j(k), \quad (19)$$

so that a weight changes only when the activation of its target concept changes [36]. Hebbian learning is simple and unsupervised, but because no error signal guides the updates, it can converge slowly and may reinforce spurious co-activations rather than correct them.

5.4 Evolutionary and Swarm Intelligence Algorithms

Evolutionary algorithms treat the weight matrix as an individual in a population and apply selection, crossover, and mutation to evolve better solutions. Genetic algorithms have been widely used to optimize FCM weights for classification and prediction tasks [37–39]. A fitness function evaluates how well a candidate weight matrix reproduces desired outputs or minimizes prediction error. Real-coded genetic algorithms represent weights as real numbers and apply arithmetic crossover and mutation operators.

PSO is a population-based search method in which each particle represents a candidate solution and moves through the search space according to its own best position and the global best position. For FCM training, PSO updates the weight matrix $W^{(\ell)}$ of particle ℓ using

$$v_{ij}^{(\ell)}(t+1) = \omega v_{ij}^{(\ell)}(t) + c_1 r_1 (w_{ij,\text{best}}^{(\ell)} - w_{ij}^{(\ell)}(t)) + c_2 r_2 (w_{ij}^{\text{gbest}} - w_{ij}^{(\ell)}(t)), \quad (20)$$

$$w_{ij}^{(\ell)}(t+1) = w_{ij}^{(\ell)}(t) + v_{ij}^{(\ell)}(t+1), \quad (21)$$

where $v_{ij}^{(\ell)}$ is the velocity of particle ℓ on weight w_{ij} , ω is an inertia factor, c_1 and c_2 are cognitive and social coefficients, $r_1, r_2 \sim \mathcal{U}(0, 1)$ are random numbers drawn uniformly, $w_{ij,\text{best}}^{(\ell)}$ is the best weight found by particle ℓ , and w_{ij}^{gbest} is the global best weight found so far [39–41]. PSO can be augmented with memetic local search to refine solutions [41].

5.5 Hybrid, Reinforcement, and Data-Driven Learning

Hybrid approaches combine Hebbian adaptation with evolutionary optimization to balance local learning and global search [49,50]. Reinforcement learning frameworks treat the FCM as an agent interacting with an environment. Rewards are associated with desirable system states, and weights are updated to maximize cumulative reward [42,43]. For example, Q-learning can adjust weights based on the temporal-difference error between predicted and actual reward signals. The data-driven end of the spectrum is broader still: convex-optimization formulations fit high-order maps directly to time series [47], neuro-fuzzy and deep FCMs integrate neural networks to learn feature representations and update rules concurrently [19,20], and self-organizing formulations that grow or prune concepts as data demand, informed by the fuzzy and probabilistic interpretations available for self-organizing networks [48], promise automatic selection of model complexity.

5.6 Comparative Analysis of Learning Algorithms

The performance of FCM learning algorithms depends on several factors, including the amount of available data, the level of complexity, and the computational budget. No single method excels in all situations; instead, practitioners should balance interpretability, complexity, and scalability when selecting an approach. Table 8 summarizes key attributes of representative learning algorithms, indicating whether supervision is required, the relative computational complexity, the ability to scale to large models, and references to key studies.

Table 8: Qualitative comparison of learning algorithms for FCMs.

Algorithm	Supervision	Complexity	Scalability	References
Hebbian	Unsupervised	Low	High	[35]
Nonlinear Hebbian	Unsupervised	Low	Medium	[36]
Genetic algorithm	Supervised/unsupervised	Medium	Medium	[37–39]
Particle swarm	Supervised/unsupervised	Medium	Medium	[39–41]
Memetic PSO	Supervised	High	Medium	[41]
Reinforcement	Reinforcement	High	Medium	[42,43]
Gradient descent	Supervised	Medium	Medium	[44,45]
Deep FCM	Supervised	High	Low	[19,20]
Self-organizing	Unsupervised	Medium	High	[48]
Hybrid evolutionary	Supervised/unsupervised	High	Medium	[49,50]

6 Variants and Extensions

Researchers have proposed numerous extensions to the basic FCM framework, each addressing a limitation of the classical model. The systematic review by Schuerkamp and Giabbanelli [1] traced this pressure precisely: classical FCMs cannot handle nonlinear relationships or time delays. They cannot fully represent uncertainty, prompting a steady stream of extensions. The same review surveyed 26 extensions proposed between 2011 and March 2023 and found that few address multiple limitations at once and that none provided code for reproducibility. To make this landscape navigable, the present section organizes the variants into three functional groups: (a) uncertainty-oriented variants, which enrich the representation of weights and activations to express imprecision (gray, interval, neutrosophic, and rough maps, with type-2 fuzzy extensions belonging to the same family); (b) temporal and dynamic variants, which enrich the reasoning process itself (high-order, random, and dynamic maps); and (c) machine-learning-integrated variants, which couple FCM inference with neural or quantum-inspired computation (deep, quantum-inspired, and chaotic maps). The subsections below follow this grouping, and Table 9 compares the variants across all three groups.

Table 9: Comparison of FCM variants and their key features.

Variant	Weight Type	Concept Type	Dynamics	Advantages
Standard FCM	Crisp $w_{ij} \in [-1, 1]$	Scalar	Static/dynamic	Simplicity, interpretability
Gray FCM	Interval $[w_{ij}^{\min}, w_{ij}^{\max}]$	Interval	Interval arithmetic	Uncertainty modeling
Interval FCM	Interval	Interval	Interval arithmetic; fixed point	Robustness to noise
Rough FCM	Crisp	Rough sets	Set-based update	Vague concept representation
Neutrosophic FCM	Triplet (T, I, F)	Triplet	Triplet update functions	Indeterminacy handling
High-order FCM	Higher tensors	Scalar	Multi-interaction	Captures high-order causal effects
Random FCM	Random variables	Scalar	Stochastic dynamics	Monte Carlo analysis
Dynamic FCM	Time-varying	Scalar	Time-dependent	Adaptation to nonstationarity
Deep FCM	Crisp	Neural layers	Deep updates	Representation learning
Quantum-inspired FCM	Complex amplitudes	Probabilistic states	Amplitude-based search	Escapes local optima; sparser maps
Chaotic FCM	Crisp	Scalar	Chaotic maps	Modeling of irregular dynamics

6.1 Gray and Interval FCMs

Gray systems theory deals with incomplete information by representing uncertain quantities as intervals. Gray FCMs allow both weights and concept activations to be gray numbers, effectively capturing uncertainty in expert knowledge and data [29]. Interval FCMs generalize this idea by associating each weight with an interval $[w_{ij}^{\min}, w_{ij}^{\max}]$, resulting in interval activation values. Update rules propagate interval arithmetic, and stability conditions require careful analysis to prevent interval explosion [30].

6.2 Neutrosophic and Rough Cognitive Maps

Neutrosophic logic extends fuzzy logic by introducing a truth–indeterminacy–falsity triad. Neutrosophic cognitive maps utilize triplets (T, I, F) to represent degrees of truth, indeterminacy, and falsity for each weight, thereby providing a richer representation of uncertainty [31]. Rough cognitive maps are based on rough set theory, which represents each concept as a pair of lower and upper approximations. Weights indicate the strength of influence between these approximations, allowing the model to handle incomplete and vague information [1].

6.3 High-Order, Random, and Dynamic Maps

High-order FCMs include interactions among groups of concepts, rather than just pairwise relations. For example, a second-order FCM introduces pair weights w_{ijl} that quantify how the joint activation of concepts C_j and C_l influences C_i , as formalized in Section 4 [32]. Random FCMs model uncertainties by treating weights as random variables drawn from specified distributions, enabling stochastic simulation of system behavior [33]. Dynamic FCMs allow weights to evolve according to adaptation rules or external inputs. For instance, a weight may decay exponentially toward a baseline value or adjust according to a feedback signal [34]. Dynamic maps can capture nonstationary phenomena and adaptive control.

6.4 Deep, Quantum-Inspired, and Chaotic FCMs

The increasing complexity of modern systems has prompted the development of advanced FCM variants that draw on deep learning, quantum-inspired computation, and the theory of chaotic dynamics. These models extend the expressive power of standard FCMs and open connections to broader areas of computational intelligence. The following subsections outline the key features, advantages, and challenges of each.

6.4.1 Deep FCMs

Deep FCMs integrate neural network architectures with FCM reasoning to automatically extract hierarchical representations from data. In a typical deep FCM, the input activation vector is first processed by one or more hidden layers of a feed-forward or recurrent neural network, and the resulting latent representation is passed through an FCM inference layer that applies the weight matrix and activation function [19]. Alternatively, the deep network may generate instance-specific weight matrices conditioned on the inputs, as in Neural-FCM variants. Some deep FCM designs for time-series forecasting use a fully connected network to map exogenous variables to latent features, an FCM layer to propagate causal influences, and partial derivatives of the output with respect to the input to preserve interpretability [20]. This architecture combines the representation-learning capabilities of deep learning with the transparent causal reasoning of FCMs, capturing nonlinear and nonmonotonic interactions among variables. On the benchmarks reported in the cited studies, deep FCMs improved classification and forecasting accuracy relative to population-based training; the magnitude of the gains varies with the dataset and baseline and is specific to each experiment [19,20].

The trade-offs are noteworthy. As neural depth increases, the causal meaning of intermediate features becomes less transparent, and the learned weight matrices may encode complex interactions that are difficult to interpret. Training deep FCMs also demands more data and computational resources than traditional FCMs. Judged explicitly against the four interpretability levels of [Section 9](#), the picture is sobering: static interpretability survives only within the FCM layer, since the encoder weights and any instance-conditioned weight generation admit no causal reading; dynamic interpretability is partial at best, because the activation trace can be inspected in the FCM layer but not through the latent representations that feed it; causal-effect interpretability is compromised whenever the effective weights are functions of the input, as the same edge can then carry different strengths for different instances; and actionable interpretability is only nominally attained, because the underlying counterfactual explanations, although technically obtainable by gradient search, operate partly in a latent space that domain experts cannot inspect. A deep FCM is therefore not interpretable by construction; whether it explains anything must be demonstrated with the explanation-evaluation metrics (fidelity, stability, compactness) discussed in [Section 9](#), and none of the deep FCM studies surveyed here reports such measurements. Recent work on Neural-FCM classifiers addresses these issues by restricting the deep component to a shallow feed-forward network and imposing sparsity constraints on the generated weight matrix; the resulting model improves classification accuracy while retaining partial interpretability and converges under cross-entropy training [\[52\]](#). Further research is needed to understand how to regularize deep FCMs to maintain causal interpretability, determine appropriate network architectures for different tasks, and develop efficient training algorithms.

6.4.2 Quantum-Inspired and Quantum FCMs

Quantum-inspired FCMs draw on concepts from quantum computation, such as superposition and interference, to represent causal influences. In these models, each edge weight is encoded as a quantum bit, with the amplitude reflecting the probability that a causal link exists and the phase encoding its sign and strength. A quantum-inspired evolutionary algorithm maintains a population of quantum chromosomes representing candidate weight matrices and applies quantum rotation gates and particle swarm operators to update them; the probabilistic representation allows the search to escape local optima and yields sparser, more accurate FCM structures than classical evolutionary approaches [\[53\]](#). The final quantum state is collapsed to produce a crisp weight matrix for inference.

A distinction is necessary here. The models described above are quantum-inspired: they borrow the mathematical formalism of superposition and probabilistic amplitudes as a search heuristic executed entirely on classical hardware, with no quantum-mechanical resources involved. Genuinely quantum FCMs, in which states or weight updates would be realized on quantum processors, remain essentially unexplored; no experimental demonstration on quantum hardware has been reported to date, and the term should not be conflated with the classical quantum-inspired algorithms surveyed here. This assessment is corroborated by the dedicated study of Kalampokas and Papakostas, who ask precisely to what extent FCMs can be quantum: their systematic review of the work at the intersection of quantum computing and FCMs offers a theoretical algorithmic assessment of how FCM representation, reasoning, and learning could map onto quantum formalisms, examines the potential improvements alongside the obstacles involved, and concludes that the integration remains largely understudied, with existing proposals confined to theory and classical simulation rather than hardware realization [\[54\]](#). With that distinction in place, quantum-inspired FCMs offer two advantages: they embed uncertainty about the very existence of causal links directly in the model, and their search operators accelerate convergence while escaping local optima. In the experiments of Ma et al., quantum-inspired FCM classifiers and ensembles built from them achieved lower error and better structure identification than classical evolutionary approaches on benchmark datasets [\[53\]](#). Nonetheless,

several challenges remain. The underlying theory is still nascent: how superposed causal relations should be interpreted, how the consistency of the update rules can be ensured, and how quantum-inspired evolutionary algorithms scale to large concept sets all require further investigation. Judged against the four interpretability levels of [Section 9](#), quantum-inspired FCMs fare acceptably at the static level, because the amplitude representation is collapsed to a crisp weight matrix whose signed edges read as usual, but the quantum-formalism search process contributes nothing explanatory in itself, and for any model that retained superposed weights, the dynamic, causal-effect, and actionable (counterfactual) levels are currently undefined because superposed causal relations have no agreed semantics. Claims that quantum inspiration improves anything beyond optimization should therefore be treated with caution. Quantum-inspired FCMs should therefore be viewed as a promising but exploratory direction. Genuinely quantum realizations remain an open research frontier rather than an established variant, a conclusion that the analysis of Kalampokas and Papakostas [54] both supports and sharpens by cataloging the representational and algorithmic questions any genuinely quantum FCM would first have to resolve.

6.4.3 Chaotic FCMs

Chaotic FCMs incorporate nonlinear chaotic maps into the activation update rules to model irregular, sensitive, and long-range dynamics [24]. Instead of applying a monotonic activation function to the weighted sum, a chaotic map, such as the logistic or tent map, is used, yielding sequences of concept activations that can exhibit deterministic chaos. Chaotic FCMs are particularly well-suited to systems whose underlying dynamics are known to be chaotic, such as ecological populations or economic cycles. By adjusting control parameters, the model can transition between periodic, quasi-periodic, and chaotic regimes, offering a range of qualitatively different behaviors. However, training chaotic FCMs is challenging because small changes in weights or activation functions can drastically alter the long-term dynamics. Noise amplification and parameter sensitivity must be carefully managed, and further research is needed to develop robust training algorithms and stability criteria for these models.

7 Stability and Convergence

Understanding stability and convergence is essential for designing reliable FCMs: a map that oscillates or drifts chaotically cannot support reproducible what-if analysis. Nápoles et al. investigated the convergence of FCMs used in pattern recognition, observing that these recurrent models propagate an activation vector through the causal network until either a fixed point is reached or a maximum number of cycles is completed, and that unstable configurations are mostly attributable to the weight matrix; they proposed a population-based learning algorithm with explicit convergence features to improve the stability of sigmoid FCM classifiers [44]. Building on these results, Nápoles et al. reviewed FCM models for pattern classification, summarizing advances and challenges in the field [55]. These insights motivate learning schemes that incorporate convergence criteria from the outset.

The formal conditions warrant careful statement. For the linear update $\mathbf{A}(k+1) = \mathbf{W}\mathbf{A}(k)$, convergence is governed by the spectral radius: $\rho(\mathbf{W}) < 1$ drives every trajectory to the origin. When an elementwise activation function f with Lipschitz constant L is applied, as in [Eq. \(1\)](#), the standard sufficient condition follows from a contraction argument: if $L\|\mathbf{W}\| < 1$ for an induced matrix norm compatible with the elementwise bound (the spectral norm $\|\mathbf{W}\|_2$ is the usual choice, with $L = \lambda/4$ for the sigmoid of [Eq. \(2\)](#) and $L = \lambda$ for the hyperbolic tangent), then the update map is a contraction, and the Banach fixed-point theorem guarantees a unique, globally attracting equilibrium. Because $\rho(\mathbf{W}) \leq \|\mathbf{W}\|$ for every induced norm, the frequently quoted condition $L\rho(\mathbf{W}) < 1$ is implied by the contraction condition but is not by itself sufficient in the nonlinear case, a subtlety that parts of the literature gloss over. Lyapunov methods offer a complementary route: taking the candidate function

$$V(\mathbf{A}) = \frac{1}{2} \|\mathbf{A} - \mathbf{A}^*\|^2, \quad (22)$$

where $\mathbf{A}^*$ is the equilibrium activation vector, the equilibrium is asymptotically stable whenever $\Delta V(\mathbf{A}) = V(\mathbf{A}(k+1)) - V(\mathbf{A}(k)) < 0$ for all $\mathbf{A} \neq \mathbf{A}^*$. For dynamic FCMs with time-varying weights, stability conditions involve the joint spectral radius of the set of weight matrices, the standard theory of which is developed in the monograph [56]. Table 10 summarizes these criteria.

Table 10: Summary of stability criteria for FCM dynamics.

Condition	Description	References
$\rho(W) < 1$	Convergence of the linear (activation-free) update	[7,9]
$L\ W\ < 1$	Contraction condition for Lipschitz activations; implies the weaker $L\rho(W) < 1$	[23]
Nonlinear contractive	Update map contractive under a general (possibly nonmonotonic) activation	[24]
Joint spectral radius	Switching or time-varying weight matrices (dynamic FCMs)	[56]

7.1 Advanced Stability Results and Open Questions

Recent research has sought to generalize the classical criteria and provide sharper bounds for the convergence of FCM dynamics. Whereas the early results bound the product of the activation function's Lipschitz constant and a norm of the weight matrix by unity, newer analyses derive sufficient conditions that relax this restriction by exploiting diagonal dominance and the distribution of weight magnitudes. A notable direction is the theoretical treatment of interval-valued FCMs. Such work formalizes the inference mechanism using interval arithmetic, shows that interval-valued FCMs reduce to conventional FCMs as interval widths shrink, and establishes convergence guarantees through spectral radius analysis and the Banach fixed-point theorem for sigmoid activations [30]. By deriving sufficient stability conditions for interval weights, this line of research addresses a gap in prior work and demonstrates the viability of FCMs under measurement uncertainty.

Another line of work provides global stability bounds for classical FCMs by analyzing the induced matrix norm of the weight matrix and proving that a unique fixed point exists when this norm falls below a threshold determined by the activation function. Conditions for the global asymptotic stability of FCMs have been derived and shown to inform parameter selection in practical applications [44,55]. Extensions employing algebraic representations based on the semi-tensor product further yield explicit formulas for the number of fixed points and limit cycles in k -valued FCMs [57]. These advances illustrate how modern tools from matrix analysis and discrete dynamical systems can sharpen the understanding of FCM stability.

Collectively, these developments indicate that stability analysis is moving beyond simple spectral radius criteria to encompass interval and stochastic weights, nonmonotonic activation functions, and multiple equilibria. They also highlight that global stability may not always be desirable: complex systems with multiple equilibrium points can exhibit meaningful multistability or oscillatory behavior. Chaotic FCMs explicitly leverage sensitive dependence on initial conditions to model irregular patterns [24]. Open questions include establishing stability criteria for gray, rough, and neutrosophic FCMs, characterizing the emergence of limit cycles and chaotic attractors in high-dimensional maps, and designing learning algorithms that balance stability with the ability to represent complex behaviors. Addressing these challenges will require interdisciplinary approaches combining spectral analysis, Lyapunov theory, contraction mapping, and bifurcation analysis.

8 Evaluation Metrics and Datasets

A survey of methods is of limited use without a clear account of how their results are judged, so this section reviews the metrics and data sources on which FCM studies rest. For classification problems, the standard metrics require only verbal definition, as they are textbook material: accuracy is the fraction of instances classified correctly; precision is the fraction of predicted positives that are truly positive; recall is the fraction of actual positives that the model recovers; and the F1-score is the harmonic mean of precision and recall. On imbalanced datasets, macro-averaged F1 is more informative than raw accuracy and should be preferred. For regression tasks, the mean squared error (MSE) is

$$\text{MSE} = \frac{1}{m} \sum_{p=1}^m (y_p - \hat{y}_p)^2, \quad (23)$$

where y_p and $\hat{y}_p$ are the observed and predicted values for pattern p . When comparing FCM learning algorithms, it is equally important to report computational complexity, convergence speed, and robustness to noise, since two algorithms with similar errors can differ sharply in these respects.

No standard benchmark repository exists for FCMs; instead, studies draw on datasets from the applications that motivated them. Table 11 lists representative examples, along with the modeling task each supports [58–61]. Many studies also construct synthetic FCMs to evaluate learning algorithms under controlled conditions, and data from time series, expert surveys, and observational studies provide further testbeds for FCM optimization. The absence of shared benchmarks is not a cosmetic gap: it is the main reason the comparative evidence summarized in Table 7 remains qualitative.

Table 11: Representative data sources used for FCM modeling and evaluation.

Data Source	Domain	Modeling Task	Reference
Process control system	Engineering	Supervisory control of tanks and valves	[35]
Language impairment cases	Medicine	Differential diagnosis	[58]
Eco-industrial park	Ecology	Environmental impact analysis	[59]
Corporate financial data	Finance	Financial distress prediction	[60]
Service supply chain	Management	Performance evaluation	[61]

9 Interpretability and Human Factors

One distinguishing feature of FCMs is their interpretability: the graph structure and weight assignments expose explicit causal pathways that domain experts can readily follow. Interpretability is not guaranteed, however; high-dimensional maps and nonlinearity can obscure causal reasoning and invite cognitive biases. To enhance transparency, Hoitsma et al. introduced a symbolic explanation module that generates human-readable accounts of the reasoning process in FCM-based models, demonstrating how FCMs can be paired with symbolic AI to produce narrative explanations [62]. A complementary route to assessing interpretability is to measure agreement between FCMs constructed by independent experts. The fuzzy Jaccard similarity index for two activation vectors $\mathbf{A}$ and $\mathbf{B}$ with nonnegative entries is defined as

$$J(\mathbf{A}, \mathbf{B}) = \frac{\sum_{i=1}^n \min(A_i, B_i)}{\sum_{i=1}^n \max(A_i, B_i)}, \quad (24)$$

where $0 \leq J \leq 1$ and larger values indicate closer agreement. Another measure is the cosine similarity between the weight matrices of two maps,

$$\cos(\theta) = \frac{\sum_{i,j} w_{ij}^{(1)} w_{ij}^{(2)}}{\sqrt{\sum_{i,j} (w_{ij}^{(1)})^2} \sqrt{\sum_{i,j} (w_{ij}^{(2)})^2}}, \quad (25)$$

where the superscripts index the two FCMs; a high value implies that the maps encode similar causal structures. Human factors such as confirmation bias and availability heuristics can distort the elicitation, and structured group techniques such as the Delphi method help mitigate these biases [63,64]. Table 12 summarizes the main human factors affecting FCM interpretability and the strategies commonly used to address them.

Table 12: Human factors affecting FCM interpretability and strategies to address them.

Aspect	Challenge	Mitigation Strategies
Expert bias	Subjective weight assignments	Group elicitation, Delphi method
Model complexity	Many concepts and cycles	Modularization, hierarchical FCMs
Communication	Hard to explain to stakeholders	Visualization, sensitivity analysis
Conflict resolution	Divergent expert opinions	Aggregation metrics (cosine, Jaccard)

A modern reading of FCM interpretability benefits from distinguishing four levels at which an FCM can explain [3,4,65], a reading consistent with the explainable-decision-making perspectives consolidated in the recent review by Sovatzidi and Iakovidis [4]. Static or topological interpretability concerns the structure itself: signed, weighted edges and centrality measures reveal which concepts drive the system before any simulation is run. Dynamic interpretability concerns the reasoning trace: because inference proceeds by iterated local updates, activation trajectories can be inspected step by step, something opaque models do not offer. Causal-effect interpretability concerns what-if analysis: clamping or perturbing a concept and propagating the change yields effect estimates whose validity, however, depends on whether the learned weights admit a causal rather than merely associational reading, a caveat that data-driven training frequently violates [65]. A second, more practical obstacle at this level has been computational: enumerating causal pathways grows combinatorially with map size. Recent advances in scalable causal-effect analysis have removed much of this barrier; in particular, the Total Causal Effect Calculation for FCMs (TCEC-FCM) algorithm efficiently computes total causal effects among concepts in large-scale FCMs by combining binary search with graph-traversal techniques, making rigorous what-if auditing feasible for maps well beyond the hand-inspection size [66]. Actionable interpretability, finally, concerns counterfactual explanations: identifying minimal changes to input concepts that flip a decision, which the graph structure of FCMs makes computationally natural. Long underexplored, this level has now received its first dedicated FCM-specific treatment: Tyrovolas, Nápoles, and Stylios derive backpropagation-based counterfactual explanations for quasi-nonlinear FCMs, exploiting the differentiability of the reasoning rule to compute minimal input changes that alter the model's decision [67]. Extending counterfactual generation to other FCM families, and evaluating such explanations with users, remain open tasks. Most published FCM applications, it should be said plainly, stop at the first level; the claim that FCMs are interpretable per se is therefore a claim about potential rather than about delivered explanation.

Interpretability also interacts with fairness. Because FCM weights propagate influence transitively, a bias encoded in one edge (for example, an expert-elicited link involving a protected attribute such as gender or ethnicity) can propagate through the causal map and contaminate downstream concepts even when the protected attribute is not directly connected to the output. Hidden confounding is equally possible: two concepts may co-activate because of an unmodeled common cause, and a learned weight between them will

then encode a spurious relation that reads, misleadingly, as causal. Fairness-aware FCM evaluation is largely an open problem; at a minimum, practitioners deploying FCMs in sensitive domains should audit paths from protected attributes to decisions, report which weights were expert-elicited vs. learned, and test the sensitivity of outputs to interventions on protected concepts.

Evaluating explanations, rather than predictions, requires its own metrics [3]. Predictive accuracy says nothing about whether the accompanying explanation is faithful or useful. For FCMs, explanation evaluation can draw on: fidelity, whether the cited causal pathway actually accounts for the model output when tested by ablation of the involved edges; stability, whether similar inputs yield similar explanations; compactness, the number of concepts and edges a user must inspect; and human-grounded measures such as comprehension accuracy and decision time in user studies. The near absence of such evaluations in published FCM applications is a substantive gap, and reporting at least fidelity and stability alongside predictive metrics would materially strengthen future empirical work. Encouragingly, systematic comparisons of explainability between FCMs and other interpretable modeling families have begun to appear [68], offering a template for the kind of head-to-head explanation evaluation the field needs.

10 Comparison with Other Modeling Frameworks

Choosing a modeling framework is rarely a matter of one paradigm dominating another; it is a matter of matching representational commitments to the problem at hand. Table 13 contrasts FCMs with Bayesian networks (BNs), artificial neural networks (ANNs), system dynamics (SD) models, and agent-based models (ABMs) along five dimensions, and the discussion that follows draws out when each choice is warranted.

Table 13: Comparison of FCMs with other modeling frameworks.

Framework	Causality	Uncertainty	Dynamics	Interpretability	Learning
FCM	Directed weighted graph	Gray/interval variants	Iterative updates	High (graph)	Hebbian, gradient, GA, PSO, RL
BN	Directed acyclic graph	Probabilities	Static or temporal BNs	Medium	Expectation-maximization, Bayesian inference
ANN	Neural network	Weights	Discrete epochs	Low	Backpropagation, gradient descent
SD	Differential equations	Parameters	Continuous time	Medium	Parameter estimation, calibration
ABM	Agents + rules	Stochastic behavior	Discrete event	High (local rules)	Simulation, calibration

The table serves as a second reading decision aid. BNs are the natural choice when calibrated probabilities matter, and the domain tolerates an acyclic structure. Still, they cannot represent the feedback loops that FCMs treat as first-class citizens. ANNs win on raw predictive accuracy when data are abundant, at the cost of the causal transparency that motivates cognitive mapping in the first place. SD models suit systems whose mechanisms are understood well enough to write differential equations with meaningful parameters. In contrast, FCMs deliberately trade mechanistic precision for the ability to quickly formalize qualitative

expert knowledge. ABMs capture heterogeneity and emergent behavior among many interacting agents, a level of granularity FCMs aggregate away. These strengths are increasingly combined rather than opposed: hybrid designs embed FCMs inside agent-based simulations to give each agent an inspectable mental model, an approach whose promise and pitfalls have been examined in detail for human–environment interactions [69]. When feedback, vagueness, and scarce data coincide, and stakeholders must be able to read the model, the FCM occupies a niche none of the alternatives fills.

10.1 Failure Cases and Boundary Conditions

A balanced comparison must also record where FCMs fail, and the literature contains well-documented failure cases that this survey would misrepresent by omission. The most fundamental critique concerns semantics. Carvalho argues that the quantities elicited as “causal weights” are, in many applications, correlational beliefs rather than mechanisms, that the iterated update rule has no explicit temporal semantics, and that simulation outcomes can be artifacts of arbitrary modeling choices, the activation function, its gain, and the number of iterations, rather than properties of the system being modeled; conclusions drawn from FCM simulations in the social sciences deserve particular caution for exactly this reason [70]. This critique does not invalidate FCMs, but it does mean that an FCM whose weights were never validated against interventions supports scenario exploration rather than causal inference.

A second, empirically documented failure mode concerns predictive performance. When labeled data are abundant, mainstream machine learning models routinely outperform classical sigmoid FCM classifiers. The FCM community itself has acknowledged the mechanism behind this gap: convergence toward a unique fixed point erases input information. Hence, accuracy had to be recovered either by interpolating toward the initial stimulus in quasi-nonlinear FCMs [25] or by abandoning bounded causal weights altogether in long-term cognitive networks [26,27], in both cases trading away part of the causal reading that motivates FCMs. In time-series forecasting, survey evidence points in the same direction: pure FCM forecasters are competitive on short, low-dimensional series but are outperformed by deep recurrent architectures on long, high-dimensional ones, which is why the strongest recent results come from hybrids rather than FCMs alone [32,71]. Reviews of methods and software reach a compatible verdict, noting that the scarcity of public implementations and benchmarks has allowed weak baselines to persist [1,72].

The boundary conditions are straightforward. FCMs earn their place when data are scarce but expert knowledge is available, when feedback loops dominate the phenomenon, and when stakeholders must be able to read and challenge the model. They are dominated when data are abundant and predictive accuracy is paramount (gradient-boosted ensembles and deep networks), when calibrated probabilistic queries are required (Bayesian networks), when governing mechanisms are known precisely enough for differential equations (system dynamics), when individual heterogeneity and emergence drive outcomes (agent-based models), and when maps grow beyond a few dozen concepts, at which point the $n(n - 1)$ elicitation burden and the unreadability of dense learned matrices erode the very interpretability that justified the choice [72]. Empirical FCM studies should accordingly report a non-FCM baseline as a matter of routine; the case studies of Section 12.8 illustrate how rarely this has been done.

11 Implementation Guidelines and Best Practices

Building an effective FCM entails several steps: (1) identify relevant concepts and variables; (2) determine the causal structure and assign initial weights through expert elicitation or data analysis; (3) select appropriate activation functions and learning algorithms; (4) simulate the FCM under different scenarios and evaluate performance; and (5) iterate to refine the map. A common objective function in classification tasks is the cross-entropy loss

$$L_{CE} = - \sum_{p=1}^m \sum_{i \in \mathcal{O}} (y_i^{(p)} \log \hat{y}_i^{(p)} + (1 - y_i^{(p)}) \log(1 - \hat{y}_i^{(p)})), \quad (26)$$

where $\mathcal{O}$ is the set of output concepts, $y_i^{(p)} \in \{0, 1\}$ denotes the ground truth for pattern p , and $\hat{y}_i^{(p)}$ the predicted probability. For regression tasks, one may use the coefficient of determination

$$R^2 = 1 - \frac{\sum_{p=1}^m (y_p - \hat{y}_p)^2}{\sum_{p=1}^m (y_p - \bar{y})^2}, \quad (27)$$

where $\bar{y}$ is the mean of the observed values. Correlation between predicted and observed concept activations can be measured with the Pearson coefficient

$$r = \frac{\sum_{p=1}^m (y_p - \bar{y})(\hat{y}_p - \bar{\hat{y}})}{\sqrt{\sum_{p=1}^m (y_p - \bar{y})^2} \sqrt{\sum_{p=1}^m (\hat{y}_p - \bar{\hat{y}})^2}}, \quad (28)$$

written here as r to keep the symbol ρ reserved for the spectral radius. Good modeling practice involves cross-validation, sensitivity analysis, and expert feedback throughout; Table 14 condenses these recommendations step by step.

Table 14: Guidelines for constructing and validating FCMs.

Step	Recommendations
Concept identification	Use domain knowledge, literature review, stakeholder interviews
Weight elicitation	Convert linguistic judgments via membership functions; use group consensus
Activation function	Choose based on desired range and smoothness; test alternatives
Learning algorithm	Match to data availability; combine Hebbian with evolutionary methods
Evaluation	Use cross-entropy, MSE, R^2 , correlation; perform sensitivity analysis
Refinement	Iterate with expert feedback; prune redundant concepts; normalize weights

11.1 Concept Selection and Data Integration

Selecting relevant concepts from raw data is a crucial step in building interpretable and effective FCMs. In data-driven settings, continuous variables must be connected to the abstract layer of concepts on which FCMs operate. Stylios and Groumpos describe an elicitation procedure in which experts first identify a set of salient variables and then group them into higher-level concepts based on semantic similarity and causal relevance [73]. Influence strengths are expressed through linguistic variables ranging from “negatively very strong” to “positively very strong” and converted into numerical weights through triangular membership functions. Experts choose the central values of the membership functions, which are aggregated across respondents using center-of-gravity defuzzification [74]. Concepts can be derived from domain knowledge, clustering of correlated variables, or principal component analysis; the latter compresses high-dimensional sensor or survey data into fewer interpretable concepts. Automated approaches use graph-theoretic metrics such as node degree, betweenness, and total influence to identify influential variables and reduce the number of concepts while preserving predictive performance [17]. Evolutionary search can select the subset of candidate concepts that best balances accuracy against interpretability, as demonstrated by structure-optimization genetic algorithms that penalize model density while minimizing prediction error [39]. Once the concepts are defined, initial weights are elicited through expert judgment or learned from data using membership functions (Section 4.2). Iterative refinement of the concept set through expert

feedback, sensitivity analysis, and cross-validation helps ensure that the FCM captures the essential causal relationships while remaining tractable. When numerical data are available, concept values are typically normalized to the unit interval and interpreted as the degree to which a concept is “activated” in a given scenario; this normalization bridges the quantitative environment and the FCM’s fuzzy conceptual layer.

11.2 Regression and Classification Tasks

FCMs can support both classification and regression tasks, depending on the nature of the output concepts. For classification, the map is usually partitioned into an input block comprising features and an output block comprising class labels [44]. An FCM classifier does not necessarily wait for a steady state; instead, several iterations are executed to propagate information from inputs to outputs, and the class with the highest activation after a fixed number of steps is selected. One representative FCM classifier uses a fully connected recurrent structure, trains the weight matrix by gradient descent with a cross-entropy loss, and interprets the intermediate iterations as a feature-transformation stage that makes observations of the same class more compact and separable [44,45]. The design decisions include how many iterations to perform (the depth), how to initialize the weights, and how to read off the classification output. Experiments demonstrate that FCM classifiers can match the accuracy of standard machine learning methods and can serve as feature preprocessors that improve downstream classifiers. Jastrzębska and Nápoles provide a comprehensive overview of FCM-based classification schemes, highlighting design choices such as network connectivity, cost function (mean squared error vs. cross-entropy), and supervision level [45].

For regression problems, the output concepts correspond to continuous variables such as system states, risk scores, or time-series values. Here, the objective is to minimize a loss function, such as mean squared error, or to maximize the coefficient of determination (R^2) between predicted and observed values. Time-series forecasting is often cast as a regression problem, in which dynamic FCMs model temporal dependencies by linking activations across successive time steps, and deep FCMs that learn complex feature representations have shown improved forecasting accuracy on multivariate sequences [19,20]. Performance evaluation relies on MSE, R^2 , and correlation, while weight interpretability can be assessed by examining how input concepts influence the continuous outputs. Cross-validation remains essential to prevent overfitting and to tune hyperparameters such as the learning rate, the number of iterations, and the regularization strength.

Dynamic FCMs, whether employed for classification or regression, leverage the model’s recurrent nature to propagate causal influences over multiple time steps. This recurrence allows the network to capture temporal patterns or complex interdependencies. Still, it also raises questions about convergence and stability that must be addressed through the stability criteria discussed in [Section 7](#).

12 Application Domains

Before surveying individual domains, three critical observations frame what follows. First, the evidential quality of the application literature is uneven: most published FCM applications validate against a single dataset or scenario constructed by the same authors, comparisons against non-FCM baselines are the exception rather than the rule, and code or data are rarely released, so independent replication is seldom possible. This echoes Schuerkamp and Giabbanelli’s finding that none of the 26 surveyed extensions provided code for reproducibility [1]. Second, what consistently works across domains is the elicitation role of FCMs, structuring expert knowledge into an inspectable causal model, and their scenario-analysis role; what consistently remains weak is out-of-sample predictive validation, where FCMs are frequently outperformed by task-specific statistical or machine-learning models unless hybridized with them. Third, read reported successes with a publication-bias caveat, since negative results with FCMs are rarely published. The domain

subsections and the case studies in [Section 12.8](#) should be read against this backdrop, and the case studies deliberately discuss methodological limitations rather than only outcomes.

With those caveats in place, the breadth of the application record remains striking. The subsections that follow tour the major areas, from travel behavior and pandemic response through engineering, energy, health care, economics, and computing, citing key studies in each. The tour is deliberately representative rather than exhaustive. Each subsection cites studies selected for methodological visibility or historical influence rather than attempting a census, and analytical depth is concentrated in the case studies of [Section 12.8](#), which examine FCM contributions relative to conventional alternatives.

12.1 Travel Behavior Modeling

Travel behavior is shaped by many interacting psychological and contextual factors, which makes it a natural target for cognitive mapping. Leon et al. first proposed using FCMs to represent travelers' mental models, combining clustering and optimization techniques to extract and aggregate individual maps [75]. Their study developed an automated approach that learned concept weights using particle swarm optimization and genetic algorithms, grouped travelers into clusters, and aggregated the resulting FCMs to characterize segment-specific behavior, demonstrating FCMs' potential to capture travel preferences, evaluate transport policies, and integrate expert and user data. Follow-up work consolidated the methodology, formalizing how individual maps are clustered and aggregated into segment-level models through a dedicated aggregation procedure [76]. Collectively, this line of research illustrates how swarm intelligence, clustering, and map aggregation combine to model heterogeneous populations of decision-makers.

In parallel, dedicated tools and methodological frameworks have been developed to support scenario analysis and experimentation. Building on the modeling and experimentation framework conceived in Leon's early work, Nápoles et al. created an FCM-based software tool for pattern classification and scenario analysis and later released *FCM Expert*, which facilitates the interactive simulation and evaluation of FCMs [77]. These software advances complement the algorithmic contributions described above and underscore the continuing maturation of FCM research and practical development.

12.2 Pandemic Modeling and Values-Based Management

Emerging crises such as pandemics call for modeling frameworks that can capture the interdependencies among epidemiological, behavioral, and socioeconomic factors. Groumpos argued that FCMs are well suited to represent the causal relationships governing disease spread and intervention outcomes; his 2021 paper on COVID-19 applied FCMs to model the pandemic and demonstrated how scenario analysis can inform health policy [78]. Beyond epidemiology, FCMs have been used to explore the spiritual and ethical dimensions of management, highlighting the role of intangible values in organizational decision-making [79]. Few methods stretch comfortably from virus transmission to organizational values, and this range illustrates how naturally FCMs absorb both quantitative and qualitative aspects of complex systems. Apostolopoulos and Groumpos later emphasized the importance of FCMs in XAI, noting that their transparent, graph-based structure addresses the growing demand for interpretable models [3].

12.3 Engineering and Control

Industrial processes were among the first proving grounds for FCMs, and engineering remains one of their most active domains. Applications include the supervision of multiagent and manufacturing systems [80], intelligent control of industrial processes [74,81], predictive maintenance of machinery [82], supply chain management [61], spatially aware decision support [83], flood resilience assessment [84], causal analysis in ports and maritime organizations [85], urban design support [86], and the supervisory modeling

of complex plants [73]. Compared with black-box models, FCMs let engineers embed expert knowledge directly and adjust weights as systems evolve. At the same time, hybrid neuro-fuzzy and reinforcement learning approaches add adaptability in highly dynamic settings [42,43].

12.4 Energy, Environment, and Sustainability

In the energy and environmental domains, FCMs help decision-makers reason about the web of relationships connecting ecological variables, policies, and outcomes. They have been used for issue identification in water resource conflicts [87], robust climate policy design [88], local planning of renewable energy systems [89], solar energy forecasting [90], coupling stakeholder cognition with agent-based models of human–environment systems [69], assessing the cumulative effects of environmental stressors on Arctic wildlife [91], ecological impact assessment of eco-industrial parks [59], and air quality early warning [92]. Researchers have further applied FCMs to crisis management and political decision-making [93], the social dimension of sustainable development [94], and agricultural decision support [95]. Interval and gray formulations prove especially valuable here, where environmental data are chronically sparse or uncertain [29,30].

12.5 Health Care and Bioinformatics

Medical decision support is one of the most active application areas of FCMs. An early landmark in the differential diagnosis of specific language impairment used a competitive map to separate overlapping developmental disorders [58]. A recent two-decade review documents how broadly the approach has spread across medical specialties [96], and a representative modern example combines myocardial perfusion imaging, clinical data, and natural-language insights in an explainable deep FCM to diagnose coronary artery disease [97]. Explainable FCMs allow clinicians to inspect the causal relationships among symptoms, biomarkers, and treatments, thereby supporting informed and defensible decisions. Gray and interval formulations, originally developed for uncertain industrial data, offer the same protection when medical knowledge is incomplete [98], while integration with machine learning enables the automatic derivation of concept relationships from electronic health records and genomic data. FCMs have also proven effective as participatory instruments in health-related research, where stakeholders co-construct maps that inform collective decision-making [99].

12.6 Economics, Finance, and Management

FCMs provide transparent models for economic systems characterized by dense feedback and persistent uncertainty. They have been used to predict corporate financial distress with interval-valued maps [60], to build goal-oriented decision support with genetically learned maps [37], to fuse high-order FCMs with recurrent networks for financial market prediction [32], and to evaluate crowdfunding decisions [100]. In management, FCMs assist with knowledge management strategy selection [63], cost analysis of prefabricated buildings [101], customer churn prediction [102], critical-success-factor analysis in service systems [103], and sustainability assessment in through-life engineering services [104]. At the societal scale, multilevel FCMs have modeled how public opinion takes shape [105]. Across these settings, the ability to fold expert intuition into a simulable causal model is what makes FCMs attractive for strategic decision-making.

12.7 Computer Science, AI, and Robotics

Within computer science, FCMs have been adopted for knowledge representation and reasoning, robotics, and AI. Applications include the selection of virtual reality development projects [106], robot navigation and path planning [46], decision-making in autonomous vehicles and drones [107], adaptive

e-learning and educational assessment [43,64,108], link prediction in social networks [109], security risk assessment [110], and knowledge-oriented causal inference from data [111]. Deep FCMs and quantum-inspired FCMs push the frontier by importing ideas from deep learning and quantum computing [19,20,51]. Software tools such as FCMpy offer user-friendly implementations of FCM construction, simulation, and optimization [28], and information-flow-based FCMs have been proposed to sharpen interpretability for human-machine interaction and control [112]. In manufacturing and industrial settings, learning FCMs have delivered explainable fault prediction, demonstrating the practical value of transparent models [113]. A growing body of work positions FCMs as a bridge between symbolic and sub-symbolic AI, arguing that they can connect symbolic reasoning with neural approaches while preserving both transparency and performance [114,115]. Table 15 presents a non-exhaustive mapping between domains and representative references.

Table 15: Representative application domains and references for FCMs.

Domain	Representative References
Manufacturing and control	[73,74,80,81]
Supply chain and logistics	[38,61,82]
Energy and sustainability	[88–90,94]
Environmental assessment	[59,87,91,92]
Health care and medicine	[58,96,97,99]
Economics and finance	[32,37,60,100]
Management and business	[63,101,102,104]
Public policy and crisis management	[12,78,93,105]
Education and training	[43,64,108]
Robotics and autonomous systems	[46,106,107]
Social networks and security	[109,110]
AI and machine learning	[19,20,51,111]

12.8 Case Studies in Depth

To complement the broad survey of domains, this subsection examines three representative case studies in greater depth. The examples span industrial process control, participatory environmental decision-making, and medical diagnosis, and together they trace the full modeling arc of concept identification to weight elicitation, learning, simulation, and evaluation.

Industrial process control. Papageorgiou, Stylios, and Groumpos demonstrated their Active Hebbian Learning (AHL) algorithm on a chemical process control problem comprising two tanks, three valves, a heating element, and paired thermometers [35]. The FCM contained eight concepts, covering the liquid level and temperature in each tank and the states of the valves and heater, with the control objective of keeping levels and temperatures within prescribed operating ranges. Experts specified the initial causal structure and weights, and the AHL procedure refined the weights through sequential concept activation. In the reported simulations, the trained map drove the process into its desired operating regions, illustrating how unsupervised refinement can adapt an elicited model without sacrificing its readability. Set against conventional alternatives, the contribution is precise: a proportional–integral–derivative (PID) or rule-based supervisory scheme could have regulated the same set-points, but it would not have exposed the causal couplings among levels, valves, and heating in a form operators can inspect and adjust; conversely, the study reports no head-to-head comparison against such a conventional controller, so the case demonstrates added transparency at an unmeasured cost in control performance.

Participatory flood resilience assessment. Mehryar and Surminski used fuzzy cognitive mapping to study flood resilience in Lowestoft, a coastal town in the United Kingdom exposed to tidal, pluvial, and fluvial flooding [84]. Twenty-three individual maps were elicited through focus groups and interviews with government, business, and community stakeholders, and then normalized and aggregated into a collective model comprising 40 concepts connected by 161 causal links. Rather than running dynamic simulations, the study analyzed the aggregated structure using network centrality measures and voting methods, and deliberately reported the variability of stakeholder opinions alongside mean weights rather than averaging disagreement away. A striking finding was that risk communication and stakeholder engagement exerted the greatest system-wide influence on resilience even though participants had ranked them near the bottom of their stated priorities, precisely the kind of insight that motivates structural FCM analysis. This is also what a conventional alternative would have missed: a survey-based priority ranking or a regression on stated importance reproduces the participants' rankings, in which risk communication sat near the bottom, whereas the centrality analysis of the aggregated map surfaced its system-wide leverage. The price is that the finding remains structural; no outcome data validate whether acting on it improves resilience.

Medical differential diagnosis. Georgopoulos et al. built a competitive FCM for the differential diagnosis of specific language impairment, a developmental disorder that is easily confused with severe dyslexia or mild autism [58]. The map comprises eighteen concepts: fifteen factor concepts encoding symptoms and causative factors, and three mutually inhibiting disorder concepts, so that iterating the sigmoid update drives the competing diagnoses apart until one dominates. Weights were derived from published clinical knowledge through linguistic variables and fuzzy aggregation. When evaluated on documented clinical cases of all three disorders, the model reproduced the correct diagnosis even when the available information was incomplete, and large ensembles of randomized inputs confirmed the hypothesized overlap among the disorders. The contribution relative to standard classifiers is equally specific: a logistic regression or decision tree of comparable accuracy would have required a labeled case base that did not exist for this rare differential, whereas the competitive map was assembled from published clinical knowledge and degrades gracefully under missing inputs. What the study does not show, because no comparison on a common dataset was run, is whether a data-trained classifier would match or exceed its accuracy wherever labeled cases are available. Table 16 summarizes the three case studies.

Table 16: Overview of the selected case studies illustrating FCM applications.

Case Study	Domain	Concepts	Variant	Weight Derivation	Reference
Process control	Engineering	8	Standard	Expert map, Active Hebbian Learning	[35]
Flood resilience	Environment	40	Participatory (aggregated)	Stakeholder elicitation, map aggregation	[84]
Differential diagnosis	Medicine	18	Competitive	Clinical literature, fuzzy aggregation	[58]

Three methodological limitations cut across these cases and reflect the application literature at large. First, none of the three studies released code or machine-readable data, so readers cannot independently reproduce the reported results and must rely on the authors' evaluations. Second, each study examines a single system, town, or clinical population, so external validity is unknown, and the reported successes

were measured under conditions chosen by the original authors. Third, in all three cases, experts fixed the concept set in advance, which keeps the model interpretable but means the evaluation cannot detect omitted-variable problems. These caveats do not diminish the value of the studies; they explain why the field still lacks the cumulative, comparable evidence base that shared benchmarks would provide. Read together with [Section 10.1](#), they also delimit what the three cases establish: in each, the FCM added elicitation transparency and structural insight that a conventional model would not have produced, but none demonstrates predictive superiority over a conventional baseline, and claims on that dimension should not be inferred from them.

13 Software Tools

Good tooling is what turns a modeling methodology into daily practice, so this section reviews the software available for building and analyzing FCMs and the features each tool provides.

FCMpy is an open-source Python library that supports the construction, simulation, and analysis of FCMs, including implementations of nonlinear Hebbian learning (NHL), AHL, and GA; it also derives weight matrices from qualitative survey data via membership functions and provides scenario-analysis routines [28]. Mental Modeler, a free web application originally developed for adaptive environmental management, enables stakeholders to draw maps and run one-step scenario comparisons collaboratively, making it a standard instrument in participatory modeling [116]. *FCM Expert*, an interactive Java tool introduced by Nápoles et al., supports scenario analysis and pattern classification with built-in learning algorithms [77], and it evolved from an earlier computational tool for the simulation and training of FCMs presented in 2015 [117]. Beyond dedicated packages, FCMs are regularly embedded within larger decision-support workflows, for example to rank knowledge management strategies [63]. As noted in [Section 8](#), the lack of shared benchmarks complicates comparisons among these tools; public datasets and open competitions would help standardize their evaluation.

[Table 17](#) summarizes the principal tools. Each row lists the platform, release year, availability, key features, and a representative reference.

Table 17: Software tools for FCM modeling and analysis. Each row summarizes the platform, release year, availability, key features, and a representative reference.

Software	Platform	Year	Availability	Features	Reference
FCMpy	Python	2022	Open source	Construction, simulation, scenario analysis, and learning (NHL, AHL, GA)	[28]
Mental Modeler	Web	2013	Free web application	Participatory model building and scenario analysis	[116]
FCM Expert	Java	2018	Freely available	Scenario analysis, pattern classification, and learning algorithms	[77]
Simulation and learning tool	Desktop	2015	From the authors	Simulation and training of FCM models	[117]

14 Cross-Cutting Synthesis

A survey that only catalogs leaves the reader to arbitrate among competing claims, so this section states explicitly what the assembled evidence supports, organized around four tensions that recur across the preceding sections and that the literature has often left unresolved.

The first tension concerns convergence. Stability theory ([Section 7](#)) treats convergence to a unique fixed point as the desirable regime, while the classification literature identifies exactly that behavior as the reason classical FCM classifiers lose accuracy: a map that forgets its input cannot discriminate [[25](#)]. These positions are not contradictory once the task is made explicit. For scenario analysis and decision support, unique attractors are a feature, because reproducible equilibria are what stakeholders interrogate; for pattern classification, they are a defect, and the quasi-nonlinear and long-term cognitive network families are best read as principled responses that relax, respectively, the reasoning rule and the weight bounds [[25–27](#)]. The practical rule remains the one stated earlier: adopt the weakest relaxation the task requires, because every step away from the classical rule weakens the causal reading of the weights.

The second tension concerns learning algorithms, where individual studies report conflicting claims. The conflict dissolves once two hidden variables are exposed: how far the initial map is from the truth and how much data there is. Hebbian methods refine near-correct expert maps and fail at structural correction; population-based search corrects structure on small and medium maps but scales poorly and is rarely validated beyond a few dozen concepts; gradient-based training wins when the pipeline is differentiable and data are comparatively plentiful; and reinforcement schemes are justified mainly in nonstationary control settings, where the evidence base is thinnest ([Section 5.2](#)). Reported head-to-head differences beyond these regimes largely reflect incommensurable datasets and protocols rather than intrinsic algorithmic merit, which is why this survey declines to declare a single best learner.

The third tension is the gap between claimed interpretability and delivered interpretability. The four-level framework of [Section 9](#) makes the claim testable, and applying it yields an uncomfortable result: most published applications operate at the static level only, explanation quality is rarely measured, and the actionable (counterfactual) level acquired its first dedicated FCM method only recently [[67](#)]. The position taken here is that interpretability must be treated as a measurable property, reported through fidelity, stability, and compactness, rather than as a birthright of the formalism, and the deep and quantum-inspired variants were evaluated in exactly those terms in this survey.

The fourth tension is proliferation rather than consolidation. More than two dozen extensions coexist; few address more than one limitation, and almost none ship code [[1,72](#)]; meanwhile, the uncertainty-oriented variants remain underused in the health and environmental domains where uncertainty quantification would matter most. The field's marginal returns now lie less in new variants than in shared benchmarks, released implementations, and validation studies that would let the existing ones compete on record. Taken together, the four resolutions locate the durable comparative advantage of FCMs in structured elicitation, inspectable feedback dynamics, and hybridization with stronger predictors, rather than in stand-alone predictive accuracy, and [Section 10.1](#) marks the boundaries beyond which that advantage does not extend.

15 Challenges and Future Research Directions

Despite considerable progress, several challenges still stand between FCMs and wider adoption. Many maps continue to rely on expert knowledge for their construction, which introduces subjectivity and limits scalability [[17,71](#)]; Nápoles et al. surveyed the theoretical and practical considerations of fuzzy cognitive modeling and identified open problems of exactly this kind, spanning interpretability, learning efficiency, and scale [[118](#)]. Data-driven techniques can relieve the dependence on experts, but they require curated datasets

and sophisticated optimization algorithms. High-dimensional systems yield dense weight matrices that are hard to interpret and expensive to update; dimensionality reduction, modularization, and hierarchical FCMs offer possible remedies [12,63]. Handling uncertainty remains a central research direction: gray, interval, neutrosophic, and rough FCMs open promising avenues, but each also raises model complexity and computational cost [29–31], and guaranteeing stability in dynamic and random FCMs demands rigorous analysis and robust learning rules [33].

Integration with modern AI presents opportunities and hazards in equal measure. Deep FCMs aim to borrow the representation learning power of deep networks while preserving the interpretability of the cognitive map [19,20]; quantum-inspired and chaotic FCMs open theoretical connections to quantum computing and dynamical systems [24,51]; and reinforcement learning and adaptive control frameworks let FCMs learn from continuous interaction with real systems [42,43]. The success of all these approaches hinges on the availability of large datasets, computational resources, and, above all, shared evaluation benchmarks. Future research should also pursue hybridization with other modeling paradigms, including Bayesian networks, system dynamics, and agent-based models. Finally, ethical considerations surrounding explainability, bias, and fairness will grow only more pressing as FCMs are deployed in sensitive domains such as medicine and public policy.

16 Conclusion

FCMs have grown from simple graph-based encodings of expert knowledge into a mature modeling framework with solid theoretical foundations and broad practical reach. This survey assembled that progress into a single reference. It traced the origins of FCMs; formalized concept activation dynamics, weight update rules, activation and membership functions, and supervised error functions; and organized the main learning algorithms into a unified taxonomy supported by a qualitative comparison of supervision, complexity, and scalability. It cataloged methods developed to capture uncertainty and richer dynamics, from gray, interval, neutrosophic, and rough maps to high-order, random, dynamic, deep, quantum-inspired, and chaotic formulations, and it summarized stability and convergence results alongside the questions that remain open. The survey also examined evaluation metrics and datasets; interpretability and expert-bias mitigation; the relationship between FCMs and alternative frameworks such as Bayesian networks, neural networks, system dynamics, and agent-based models; practical guidelines for building and tuning models; applications across many domains with illustrative case studies; and the current software ecosystem. Taken together, these threads show that FCMs' lasting strength lies in pairing interpretable causal structure with quantitative dynamics. The most promising path forward couples interpretability with the representational power of modern machine learning, large-scale data, and XAI, thereby extending the modeling capabilities of FCMs without sacrificing the transparency that motivated their development.

Acknowledgement: None.

Funding Statement: The author received no specific funding for this study.

Availability of Data and Materials: Not applicable.

Ethics Approval: Not applicable. This article is a literature survey and does not contain any studies with human participants or animals performed by the author.

Conflicts of Interest: The author declares no conflicts of interest.

References

- Schuerkamp R, Giabbanelli PJ. Extensions of fuzzy cognitive maps: a systematic review. *ACM Comput Surv.* 2023;56(2):1–36. doi:10.1145/3610771.
- Nápoles G, Jastrzębska A, Grau I, Salgueiro Y, Leon M. A review on fuzzy cognitive mapping: recent advances and algorithms. *Big Data Cogn Comput.* 2026;10(1):22. doi:10.3390/bdcc10010022.
- Apostolopoulos ID, Groumpos PP. Fuzzy cognitive maps: their role in explainable artificial intelligence. *Appl Sci.* 2023;13(6):3412. doi:10.3390/app13063412.
- Sovatzidi G, Iakovidis DK. A review of fuzzy cognitive maps and their perspectives for explainable decision making. *Neurocomputing.* 2026;697:134038. doi:10.1016/j.neucom.2026.134038.
- Axelrod R. *Structure of decision: the cognitive maps of political elites.* Princeton, NJ, USA: Princeton University Press; 1976.
- Zadeh LA. Fuzzy sets. *Inf Control.* 1965;8(3):338–53. doi:10.1016/S0019-9958(65)90241-X.
- Kosko B. Fuzzy cognitive maps. *Int J Man-Mach Stud.* 1986;24(1):65–75. doi:10.1016/S0020-7373(86)80040-2.
- Kosko B. Adaptive inference in fuzzy knowledge networks. In: *Readings in fuzzy sets for intelligent systems.* Amsterdam, The Netherlands: Elsevier; 1993. p. 888–91. doi:10.1016/B978-1-4832-1450-4.50093-6.
- Styblinski MA, Meyer BD. Signal flow graphs vs fuzzy cognitive maps in application to qualitative circuit analysis. *Int J Man-Mach Stud.* 1991;35(2):175–86. doi:10.1016/S0020-7373(05)80147-6.
- Taber R. Knowledge processing with fuzzy cognitive maps. *Expert Syst Appl.* 1991;2(1):83–7. doi:10.1016/0957-4174(91)90136-3.
- Dickerson JA, Kosko B. Virtual worlds as fuzzy cognitive maps. *Presence Teleoperators Virtual Environ.* 1994;3(2):173–89. doi:10.1162/pres.1994.3.2.173.
- Osoba OA, Kosko B. Fuzzy cognitive maps of public support for insurgency and terrorism. *J Def Model Simul: Appl Methodol Technol.* 2017;14(1):17–32. doi:10.1177/1548512916680779.
- Groumpos PP. Fuzzy cognitive maps: basic theories and their application to complex systems. In: *Fuzzy cognitive maps.* Berlin/Heidelberg, Germany: Springer; 2010. p. 1–22. doi:10.1007/978-3-642-03220-2_1.
- Papageorgiou EI, Salmeron JL. A review of fuzzy cognitive maps research during the last decade. *IEEE Trans Fuzzy Syst.* 2013;21(1):66–79. doi:10.1109/TFUZZ.2012.2201727.
- Papageorgiou EI. Learning algorithms for fuzzy cognitive maps—a review study. *IEEE Trans Syst Man Cybern C.* 2012;42(2):150–63. doi:10.1109/TSMCC.2011.2138694.
- Karatzinis GD, Boutalis YS. A review study of fuzzy cognitive maps in engineering: applications, insights, and future directions. *Eng.* 2025;6(2):37. doi:10.3390/eng6020037.
- Bakhtavar E, Valipour M, Yousefi S, Sadiq R, Hewage K. Fuzzy cognitive maps in systems risk analysis: a comprehensive review. *Complex Intell Syst.* 2021;7:621–37. doi:10.1007/s40747-020-00228-2.
- Feng G, Zhang L, Yang J, Lu W. Long-term prediction of time series using fuzzy cognitive maps. *Eng Appl Artif Intell.* 2021;102:104274. doi:10.1016/j.engappai.2021.104274.
- Qin D, Peng Z, Wu L. Deep attention fuzzy cognitive maps for interpretable multivariate time series prediction. *Knowl Based Syst.* 2023;275:110700. doi:10.1016/j.knosys.2023.110700.
- Borisov V, Luferov V. Neuro-fuzzy cognitive temporal models for predicting multidimensional time series with fuzzy trends. *Comput Sist.* 2020;24(3):1165–77. doi:10.13053/cys-24-3-3477.
- Özesmi U, Özesmi SL. Ecological models based on people's knowledge: a multi-step fuzzy cognitive mapping approach. *Ecol Model.* 2004;176(1–2):43–64. doi:10.1016/j.ecolmodel.2003.10.027.
- MacMillan GA, Badry NA, Sarmiento I, Grant E, Hickey GM, Humphries MM. Cree knowledge, fuzzy cognitive maps, and the social-ecology of moose habitat quality under an adapted forestry regime. *Ecol Soc.* 2024;29(4):34. doi:10.5751/ES-15508-290434.
- Stylios C, Groumpos PP. Fuzzy cognitive maps: a soft computing technique for intelligent control. In: *Proceedings of the 2000 IEEE International Symposium on Intelligent Control. Held Jointly with the 8th IEEE Mediterranean Conference on Control and Automation; 2000 Jul 19; Patras, Greece.* p. 97–102. doi:10.1109/ISIC.2000.882906.
- Song HJ, Miao CY, Shen ZQ, Wuyts R, D'Hondt M, Catthoor F. Design of fuzzy cognitive maps using neural networks for predicting chaotic time series. *Neural Netw.* 2010;23(10):1264–75. doi:10.1016/j.neunet.2010.08.003.

25. Nápoles G, Salmeron JL, Salgueiro Y. Inverse simulation learning of quasi-nonlinear fuzzy cognitive maps. *Neurocomputing*. 2025;650:130864. doi:10.1016/j.neucom.2025.130864.
26. Nápoles G, Grau I, Jastrzębska A, Salgueiro Y. Long short-term cognitive networks. *Neural Comput Appl*. 2022;34(19):16959–71. doi:10.1007/s00521-022-07348-5.
27. Nápoles G, Salgueiro Y, Grau I, Espinosa ML. Recurrence-aware long-term cognitive network for explainable pattern classification. *IEEE Trans Cybern*. 2023;53(10):6083–94. doi:10.1109/TCYB.2022.3165104.
28. Mkhitarian S, Giabbanelli P, Wozniak MK, Nápoles G, De Vries N, Crutzen R. FCMpy: a python module for constructing and analyzing fuzzy cognitive maps. *PeerJ Comput Sci*. 2022;8:e1078. doi:10.7717/peerj-cs.1078.
29. Salmeron JL, Gutierrez E. Fuzzy grey cognitive maps in reliability engineering. *Appl Soft Comput*. 2012;12(12):3818–24. doi:10.1016/j.asoc.2012.02.003.
30. Zhang J, Tao Z, Liu J, Liu X, Chen H. A hybrid interval-valued time series prediction model incorporating intuitionistic fuzzy cognitive map and fuzzy neural network. *J Forecast*. 2025;44(1):93–111. doi:10.1002/for.3181.
31. Al-subhi SH, Papageorgiou EI, Pérez PP, Mahdi GSS, Acuña LA. Triangular neutrosophic cognitive map for multistage sequential decision-making problems. *Int J Fuzzy Syst*. 2021;23(3):657–79. doi:10.1007/s40815-020-01014-5.
32. Zhang Y, Song Y, Wei G. An information fusion model based on error-dependent high-order fuzzy cognitive maps and long short-term memory network for financial market prediction. *Int J Info Tech Dec Mak*. 2026;25(7):2369–97. doi:10.1142/S0219622026500045.
33. Aguilar J. Adaptive random fuzzy cognitive maps. In: *Advances in Artificial Intelligence—IBERAMIA 2002*. Berlin/Heidelberg, Germany: Springer; 2002. p. 402–10. doi:10.1007/3-540-36131-6_41.
34. Amirkhani A, Shirzadeh M, Kumbasar T, Mashadi B. A framework for designing cognitive trajectory controllers using genetically evolved interval type-2 fuzzy cognitive maps. *Int J Intell Syst*. 2022;37(1):305–35. doi:10.1002/int.22626.
35. Papageorgiou EI, Stylios CD, Groumpos PP. Active Hebbian learning algorithm to train fuzzy cognitive maps. *Int J Approx Reason*. 2004;37(3):219–49. doi:10.1016/j.ijar.2004.01.001.
36. Martin N, Priya R. Exploration and comparative investigation on neutrosophic cognitive maps non-linear Hebbian learning algorithm in predictive model. In: *AIP Conference Proceedings*. Kattankulathur, India: AIP Publishing; 2022. p. 200023. doi:10.1063/5.0108500.
37. Khan MS, Khor S, Chong A. Fuzzy cognitive maps with genetic algorithm for goal-oriented decision support. *Int J Uncertain Fuzziness Knowl-Based Syst*. 2004;12(supp02):31–42. doi:10.1142/S0218488504003028.
38. Deng T. Optimization of residential landscape design and supply chain system using intelligent fuzzy cognitive map and genetic algorithm. *Comput Intell Neurosci*. 2022;2022(1):6321101. doi:10.1155/2022/6321101.
39. Poczęta K, Yastrebov A, Papageorgiou E. Learning fuzzy cognitive maps using structure optimization genetic algorithm. In: *Proceedings of the 2015 Federated Conference on Computer Science and Information Systems*; 2015 Sep 13–16; Lodz, Poland. p. 547–54. doi:10.15439/2015F296.
40. Oikonomou P, Papageorgiou EI. Particle swarm optimization approach for fuzzy cognitive maps applied to autism classification. In: *Artificial intelligence applications and innovations*. Berlin/Heidelberg, Germany: Springer; 2013. p. 516–26. doi:10.1007/978-3-642-41142-7_52.
41. Acampora G, Pedrycz W, Vitiello A. A competent memetic algorithm for learning fuzzy cognitive maps. *IEEE Trans Fuzzy Syst*. 2015;23(6):2397–411. doi:10.1109/TFUZZ.2015.2426311.
42. Sutton RS, Barto AG. *Reinforcement learning: an introduction*. 2nd ed. Cambridge, MA, USA: MIT Press; 2018.
43. Chrysafiadi K, Virvou M. A knowledge representation approach using fuzzy cognitive maps for better navigation support in an adaptive learning system. *SpringerPlus*. 2013;2(1):81. doi:10.1186/2193-1801-2-81.
44. Nápoles G, Papageorgiou E, Bello R, Vanhoof K. Learning and convergence of fuzzy cognitive maps used in pattern recognition. *Neural Process Lett*. 2017;45(2):431–44. doi:10.1007/s11063-016-9534-x.
45. Jastrzębska A, Nápoles G. Introduction to fuzzy cognitive map-based classification. In: *Fuzzy cognitive maps*. Cham, Switzerland: Springer Nature; 2024. p. 165–92. doi:10.1007/978-3-031-48963-1_9.

46. Joshi MM, Zaveri MA. Neuro-fuzzy based autonomous mobile robot navigation system. In: Proceedings of the 2010 11th International Conference on Control Automation Robotics & Vision; 2010 Dec 7–10; Singapore. p. 384–9. doi:10.1109/ICARCV.2010.5707354.
47. Shan D, Wang L, Lu W, Chen J. Convex optimization based high-order fuzzy cognitive map modeling and its application in time series predicting. *IEEE Access*. 2024;12:12683–98. doi:10.1109/ACCESS.2024.3355194.
48. Giordano L, Gliozzi V, Theseider Dupré D. A conditional, a fuzzy and a probabilistic interpretation of self-organizing maps. *J Log Comput*. 2022;32(2):178–205. doi:10.1093/logcom/exab082.
49. Mls K, Cimler R, Vaščák J, Puheim M. Interactive evolutionary optimization of fuzzy cognitive maps. *Neurocomputing*. 2017;232:58–68. doi:10.1016/j.neucom.2016.10.068.
50. Nápoles G, Jastrzębska A, Mosquera C, Vanhoof K, Homenda W. Deterministic learning of hybrid fuzzy cognitive maps and network reduction approaches. *Neural Netw*. 2020;124:258–68. doi:10.1016/j.neunet.2020.01.019.
51. Kolahdoozi M, Amirkhani A, Shojaefard MH, Abraham A. A novel quantum inspired algorithm for sparse fuzzy cognitive maps learning. *Appl Intell*. 2019;49(10):3652–67. doi:10.1007/s10489-019-01476-7.
52. Tziolas T, Papageorgiou K, Apostolopoulos I, Papageorgiou E. Neural-FCM: a deep learning approach for weight matrix optimization in fuzzy cognitive map classifiers. *Appl Intell*. 2025;55(13):948. doi:10.1007/s10489-025-06795-6.
53. Ma N, Fujita H, Zhai Y, Wang S. Ensembles of fuzzy cognitive map classifiers based on quantum computation. *Acta Polytechnica Hung*. 2015;12(4):7–26. doi:10.12700/APH.12.4.2015.4.1.
54. Kalampokas T, Papakostas GA. To what extent can fuzzy cognitive maps be quantum? *IEEE Nanotechnol Mag*. 2025;19(5):12–23. doi:10.1109/MNANO.2025.3585997.
55. Nápoles G, Leon M, Grau I, Vanhoof K, Bello R. Fuzzy cognitive maps based models for pattern classification: advances and challenges. In: *Soft computing based optimization and decision models*. Cham, Switzerland: Springer International Publishing; 2017. p. 83–98. doi:10.1007/978-3-319-64286-4_5.
56. Jungers R. The joint spectral radius: theory and applications. In: *Lecture notes in control and information sciences*. Vol. 385. Berlin/Heidelberg, Germany: Springer; 2009. doi:10.1007/978-3-540-95980-9.
57. Luo C, Song X, Zheng Y. Algebraic dynamics of k-valued fuzzy cognitive maps and its stabilization. *Knowl Based Syst*. 2020;209:106424. doi:10.1016/j.knosys.2020.106424.
58. Georgopoulos VC, Malandraki GA, Stylios CD. A fuzzy cognitive map approach to differential diagnosis of specific language impairment. *Artif Intell Med*. 2003;29(3):261–78. doi:10.1016/S0933-3657(02)00076-3.
59. Fons S, Achari G, Ross T. A fuzzy cognitive mapping analysis of the impacts of an eco-industrial park. *J Intell Fuzzy Syst: Appl Eng Technol*. 2004;15(2):75–88. doi:10.3233/IFS-2004-00219.
60. Hajek P, Prochazka O. Interval-valued fuzzy cognitive maps with genetic learning for predicting corporate financial distress. *Filomat*. 2018;32(5):1657–62. doi:10.2298/FIL1805657H.
61. Rezaei Pandari A, Azar A. A fuzzy cognitive mapping model for service supply chains performance. *Meas Bus Excell*. 2017;21(4):388–404. doi:10.1108/MBE-05-2016-0026.
62. Hoitsma F, Knobien A, Leon M, Nápoles G. Symbolic explanation module for fuzzy cognitive map-based reasoning models. In: *Artificial Intelligence XXXVII*. Cham, Switzerland: Springer International Publishing; 2020. p. 21–34. doi:10.1007/978-3-030-63799-6_2.
63. Yadegari M, Mohammadi S. A fuzzy cognitive map approach to selecting the best knowledge management strategy. *VINE J Inf Knowl Manag Syst*. 2025;55(6):1369–93. doi:10.1108/VJIKMS-07-2023-0186.
64. Barón HB, Crespo RG, Pascual Espada J, Martínez OS. Assessment of learning in environments interactive through fuzzy cognitive maps. *Soft Comput*. 2015;19(4):1037–50. doi:10.1007/s00500-014-1313-x.
65. Nápoles G, Grau I, Concepción L, Salgueiro Y, Papageorgiou EI. On the interpretability of fuzzy cognitive maps. *Knowl Based Syst*. 2023;281:111078. doi:10.1016/j.knosys.2023.111078.
66. Tyrovolas M, Kallimanis ND, Stylios C. Advancing explainable AI with causal analysis in large-scale fuzzy cognitive maps. *arXiv:2405.09190*. 2024. doi:10.48550/arXiv.2405.09190.
67. Tyrovolas M, Nápoles G, Stylios C. Backpropagation-based counterfactual explanations for quasi-nonlinear fuzzy cognitive maps. *IEEE Trans Syst Man Cybern: Syst*. 2026;56(3):1706–20. doi:10.1109/TSMC.2025.3647125.

68. Benito D, Quintero C, Aguilar J, Ramírez JM, Fernández-Anta A. Explainability analysis: an in-depth comparison between fuzzy cognitive maps and LAMDA. *Appl Soft Comput.* 2024;164:111940. doi:10.1016/j.asoc.2024.111940.
69. Giabbanelli PJ, Gray SA, Aminpour P. Combining fuzzy cognitive maps with agent-based modeling: frameworks and pitfalls of a powerful hybrid modeling approach to understand human-environment interactions. *Environ Model Softw.* 2017;95:320–5. doi:10.1016/j.envsoft.2017.06.040.
70. Carvalho JP. On the semantics and the use of fuzzy cognitive maps and dynamic cognitive maps in social sciences. *Fuzzy Sets Syst.* 2013;214:6–19. doi:10.1016/j.fss.2011.12.009.
71. Orang O, de Lima e Silva PC, Guimarães FG. Time series forecasting using fuzzy cognitive maps: a survey. *Artif Intell Rev.* 2023;56(8):7733–94. doi:10.1007/s10462-022-10319-w.
72. Felix G, Nápoles G, Falcon R, Froelich W, Vanhoof K, Bello R. A review on methods and software for fuzzy cognitive maps. *Artif Intell Rev.* 2019;52(3):1707–37. doi:10.1007/s10462-017-9575-1.
73. Stylios CD, Groumpos PP. Modeling complex systems using fuzzy cognitive maps. *IEEE Trans Syst Man Cybern A Syst Hum.* 2004;34(1):155–62. doi:10.1109/TSMCA.2003.818878.
74. de Souza LB, Prieto Soares P, Mendonca M, Mourhir A, Papageorgiou EI. Fuzzy cognitive maps and fuzzy logic applied in industrial processes control. In: *Proceedings of the 2018 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE)*; 2018 July 8–13; Rio de Janeiro, Brazil. p. 1–8. doi:10.1109/FUZZ-IEEE.2018.8491590.
75. Leon M, Nápoles G, Bello R, Mkrtchyan L, Depaire B, Vanhoof K. Tackling travel behaviour: an approach based on fuzzy cognitive maps. *Int J Comput Intell Syst.* 2013;6(6):1012–39. doi:10.1080/18756891.2013.816025.
76. Leon M. Aggregating procedure for fuzzy cognitive maps. *Int FLAIRS Conf Proc.* 2023;36. doi:10.32473/flairs.36.133082.
77. Nápoles G, Leon M, Grau I, Vanhoof K. FCM Expert: software tool for scenario analysis and pattern classification based on fuzzy cognitive maps. *Int J Artif Intell Tools.* 2018;27(7):1860010. doi:10.1142/S0218213018600102.
78. Groumpos PP. Why modelling the COVID-19 pandemic using fuzzy cognitive maps (FCM)? *IFAC-PapersOnLine.* 2021;54(13):305–10. doi:10.1016/j.ifacol.2021.10.464.
79. Groumpos PP. Fuzzy cognitive maps and spiritual modeling of management systems: issues and challenges. In: *Business intelligence and modelling*. Cham, Switzerland: Springer International Publishing; 2021. p. 65–83. doi:10.1007/978-3-030-57065-1_6.
80. Aguilar J. Dynamic fuzzy cognitive maps for the supervision of multiagent systems. In: *Fuzzy cognitive maps*. Berlin/Heidelberg, Germany: Springer; 2010. p. 307–24. doi:10.1007/978-3-642-03220-2_13.
81. Longo F, Padovano A, Nicoletti L, Fusto C, Elbasheer M, Diaz R. Fuzzy cognitive map-based knowledge representation of hazardous industrial operations. *Procedia Comput Sci.* 2021;180:1042–8. doi:10.1016/j.procs.2021.01.351.
82. Mazzuto G, Antomarioni S, Ciarapica FE, Bevilacqua M. Health indicator for predictive maintenance based on fuzzy cognitive maps, Grey Wolf, and K-nearest neighbors algorithms. *Math Probl Eng.* 2021;2021(1):1–21. doi:10.1155/2021/8832011.
83. Liu ZQ, Satur R. Contextual fuzzy cognitive map for decision support in geographic information systems. *IEEE Trans Fuzzy Syst.* 1999;7(5):495–507. doi:10.1109/91.797975.
84. Mehryar S, Surminski S. Investigating flood resilience perceptions and supporting collective decision-making through fuzzy cognitive mapping. *Sci Total Environ.* 2022;837:155854. doi:10.1016/j.scitotenv.2022.155854.
85. Ghafarzadeh A, Memarzadeh Tehran G, Hamidi N, Mohammadi N. Application of fuzzy cognitive map to design the causal structure and analyze the factors affecting good governance in the ports and maritime organization. *Int J Coast Offshore Environ Eng.* 2021;5(2):34–46. doi:10.52547/ijcoe.5.2.34.
86. Xirogiannis G, Stefanou J, Glykas M. A fuzzy cognitive map approach to support urban design. *Expert Syst Appl.* 2004;26(2):257–68. doi:10.1016/S0957-4174(03)00140-4.
87. Giordano R, Passarella G, Uricchio VF, Vurro M. Fuzzy cognitive maps for issue identification in a water resources conflict resolution system. *Phys Chem Earth Pt A/B/C.* 2005;30(6–7):463–9. doi:10.1016/j.pce.2005.06.012.
88. Nikas A, Doukas H. Developing robust climate policies: a fuzzy cognitive map approach. In: *Robustness analysis in decision aiding, optimization, and analytics*. Cham, Switzerland: Springer International Publishing; 2016. p. 239–63. doi:10.1007/978-3-319-33121-8_11.

89. Kyriakarakos G, Patlitzianas K, Damasiotis M, Papastefanakis D. A fuzzy cognitive maps decision support system for renewables local planning. *Renew Sustain Energy Rev.* 2014;39:209–22. doi:10.1016/j.rser.2014.07.009.
90. Orang O, Silva R, de Lima e Silva PC, Guimarães FG. Solar energy forecasting with fuzzy time series using high-order fuzzy cognitive maps. In: *Proceedings of the 2020 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE)*; 2020 Jul 19–24; Glasgow, UK. p. 1–8. doi:10.1109/FUZZ48607.2020.9177767.
91. Rooney RC, Daniel J, Mallory M, Hedd A, Ives J, Gilchrist G, et al. Fuzzy cognitive mapping as a tool to assess the relative cumulative effects of environmental stressors on an Arctic seabird population to identify conservation action and research priorities. *Ecol Solut Evid.* 2023;4(2):e12241. doi:10.1002/2688-8319.12241.
92. Gu X, Li H, Fan H. Spatiotemporal hybrid air pollution early warning system of urban agglomeration based on adaptive feature extraction and hesitant fuzzy cognitive maps. *Systems.* 2023;11(6):286. doi:10.3390/systems11060286.
93. Andreou AS, Mateou NH, Zombanakis GA. Soft computing for crisis management and political decision making: the use of genetically evolved fuzzy cognitive maps. *Soft Comput.* 2005;9(3):194–210. doi:10.1007/s00500-004-0344-0.
94. Nasirzadeh F, Ghayoumian M, Khanzadi M, Rostamnezhad Cherati M. Modelling the social dimension of sustainable development using fuzzy cognitive maps. *Int J Constr Manag.* 2020;20(3):223–36. doi:10.1080/15623599.2018.1484847.
95. Consentino F, Peri I, Litrico M, Spina D, Vindigni G. Mapping young farmers' choice to pursue geographical indication in a rural context: application of fuzzy cognitive map. *Agric Food Econ.* 2024;12(1):44. doi:10.1186/s40100-024-00340-8.
96. Apostolopoulos ID, Papandrianos NI, Papathanasiou ND, Papageorgiou EI. Fuzzy cognitive map applications in medicine over the last two decades: a review study. *Bioengineering.* 2024;11(2):139. doi:10.3390/bioengineering11020139.
97. Feleki A, Apostolopoulos ID, Moustakidis S, Papageorgiou EI, Papathanasiou N, Apostolopoulos D, et al. Explainable deep fuzzy cognitive map diagnosis of coronary artery disease: integrating myocardial perfusion imaging, clinical data, and natural language insights. *Appl Sci.* 2023;13(21):11953. doi:10.3390/app132111953.
98. Chen J, Gao X, Rong J. Enhance the uncertainty modeling ability of fuzzy grey cognitive maps by general grey number. *IEEE Access.* 2020;8:163844–56. doi:10.1109/ACCESS.2020.3021721.
99. Sarmiento I, Cockcroft A, Dion A, Belaid L, Silver H, Pizarro K, et al. Fuzzy cognitive mapping in participatory research and decision making: a practice review. *Arch Public Health.* 2024;82(1):76. doi:10.1186/s13690-024-01303-7.
100. Borrero-Domínguez C, Escobar-Rodríguez T. Decision support systems in crowdfunding: a fuzzy cognitive maps (FCM) approach. *Decis Support Syst.* 2023;173:114000. doi:10.1016/j.dss.2023.114000.
101. Luo L, Wu X, Hong J, Wu G. Fuzzy cognitive map-enabled approach for investigating the relationship between influencing factors and prefabricated building cost considering dynamic interactions. *J Constr Eng Manage.* 2022;148(9):04022081. doi:10.1061/(ASCE)CO.1943-7862.0002336.
102. Sivasankar E, Vijaya J. Hybrid PPFCM-ANN model: an efficient system for customer churn prediction through probabilistic possibilistic fuzzy clustering and artificial neural network. *Neural Comput Appl.* 2019;31(11):7181–200. doi:10.1007/s00521-018-3548-4.
103. Shabandarzadeh H, Jamkhaneh HB, Haghighat F, Khajeh F. Analysis of critical success factors in robust service systems through fuzzy cognitive map. *Int J Serv Econ Manag.* 2025;16(1):1–20. doi:10.1504/IJSEM.2025.144573.
104. Farsi M, Erkoyuncu JA. Toward an integrated sustainability assessment in through-life engineering services. *SSRN Electron J.* 2020. doi:10.2139/ssrn.3718041.
105. Sánchez H, Aguilar J, Terán O, Gutiérrez de Mesa J. Modeling the process of shaping the public opinion through multilevel fuzzy cognitive maps. *Appl Soft Comput.* 2019;85:105756. doi:10.1016/j.asoc.2019.105756.
106. Huang SF. Using linguistic VIKOR and fuzzy cognitive maps to select virtual reality games development project. *Mathematics.* 2021;9(11):1253. doi:10.3390/math9111253.

107. Singh R, Nishad DK, Khalid S, Chaudhary A. A review of the application of fuzzy mathematical algorithm-based approach in autonomous vehicles and drones. *Int J Intell Robot Appl.* 2024;9(1):344–64. doi:10.1007/s41315-024-00385-4.
108. Teh CS, Kudus AR. The application of fuzzy cognitive mapping in education: trend and potential. *TEM J.* 2024;13(2):976–91. doi:10.18421/TEM132-13.
109. Sharma U, Kandwal S, Khatri SK. A link prediction in social networks: a fuzzy cognitive map approach. In: *Ambient communications and computer systems*. Singapore: Springer; 2018. p. 463–73. doi:10.1007/978-981-10-7386-1_40.
110. Szwed P, Skrzyński P. A new lightweight method for security risk assessment based on fuzzy cognitive maps. *Int J Appl Math Comput Sci.* 2014;24(1):213–25. doi:10.2478/amcs-2014-0016.
111. Sharma A, Tselykh A, Podoplelova E, Tselykh A. Knowledge-oriented methodologies for causal inference relations using fuzzy cognitive maps: a systematic review. *Comput Ind Eng.* 2022;171:108500. doi:10.1016/j.cie.2022.108500.
112. Tyrovolas M, Liang XS, Stylios C. Information flow-based fuzzy cognitive maps with enhanced interpretability. *Granul Comput.* 2023;8(6):2021–38. doi:10.1007/s41066-023-00417-7.
113. Mansouri T, Vadera S. Explainable fault prediction using learning fuzzy cognitive maps. *Expert Syst.* 2023;40(8):e13316. doi:10.1111/exsy.13316.
114. Leon M. Fuzzy cognitive maps as a bridge between symbolic and sub-symbolic artificial intelligence. *Int J Cybern Inform.* 2024;13(4):57–75. doi:10.5121/ijci.2024.130406.
115. Leon M. Fuzzy cognitive maps bridging transparency and performance in hybrid AI systems. *Int J Soft Comput.* 2024;15(3):17–37. doi:10.5121/ijsc.2024.15302.
116. Gray SA, Gray S, Cox LJ, Henly-Shepard S. Mental modeler: a fuzzy-logic cognitive mapping modeling tool for adaptive environmental management. In: *Proceedings of the 2013 46th Hawaii International Conference on System Sciences*; 2013 Jan 7–10; Wailea, HI, USA. p. 965–73. doi:10.1109/HICSS.2013.399.
117. Nápoles G, Grau I, Bello R, Leon M, Vanhoof K, Papageorgiou E. A computational tool for simulation and learning of fuzzy cognitive maps. In: *Proceedings of the 2015 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE)*; 2015 Aug 2–5; Istanbul, Turkey. p. 1–8. doi:10.1109/FUZZ-IEEE.2015.7337859.
118. Nápoles G, Salmeron JL, Froelich W, Falcon R, Leon M, Vanhoenshoven F, et al. Fuzzy cognitive modeling: theoretical and practical considerations. In: *Intelligent Decision Technologies 2019*. Singapore: Springer; 2019. p. 77–87. doi:10.1007/978-981-13-8311-3_7.