



Fuzzy Cognitive Mapping

Abstract This chapter introduces Fuzzy Cognitive Mapping (FCM), a method for developing and analysing ‘semi-quantitative’ (i.e. using and producing indicative rather than predictive numerical values) causal models. We explain in simple language what an FCM map is made up of and the two main ways in which analysis is done. We go into some detail on how to do it yourself and provide reflections on common issues and tricks of the trade. We also discuss its roots and debates in the field since. Finally, we provide some advice and resources for getting started yourself.

Keywords Fuzzy Cognitive Mapping • Cognitive mapping • Mental models

Fuzzy Cognitive Mapping (FCM) holds a special place in our hearts, being one of the first system mapping methods we used together. We found the intuitive nature of FCM map building, and the offer of ‘quick and dirty’ exploration of a system’s dynamics, appealing. Over time, our practice, particularly in how we analyse system maps, has shifted in response to some of the drawbacks we felt the method had. It is hard for us to write this chapter without that history in mind. We worried about this causing bias, but our research for this chapter has led us to understand that the same debates we had about what was appropriate to do in analysing FCM maps have been at the centre of FCM debates more widely. To reflect this, we present both the schools of thought in how to analyse FCMs.

In the rest of this chapter, we describe what FCM is in a comprehensive but as simple as we could manage way. We describe any variety in terminology or practice we have seen, and we describe in detail how to do FCM. After the ‘how to’ section we discuss common issues and tricks of the trade. We then reflect on what the method is good and bad at, before outlining the history of FCM. Finally, we outline some practical next steps for readers wanting to use it themselves.

WHAT IS FUZZY COGNITIVE MAPPING?

FCM involves building a model of a system made up of boxes and connections. The boxes, normally referred to as ‘factors’ or ‘concepts’, can represent anything which is expressed as a variable (i.e. it makes sense to think of it going up or down). Importantly, it does not have to be quantifiable or to have data behind it which means it can be used to capture knowledge without empirical data. The connections (i.e. arrows) are often referred to as ‘edges’, and they represent causal links between factors. They do not represent a vague notion of ‘something is going on here’ but are meant to show direct causal influence. The exploration of how these causal influences propagate through a system when it is subject to change or intervention is at the heart of what an FCM is intended to do. There are, however, two main approaches to this in common use within the community, and it is often not explicitly stated which approach is being used.

The first approach is most related to the origins of FCM and retains the same mathematical formulation. In this approach, which we will call ‘causal’ (following Helfgott et al., 2015), the strength of links between factors represents how certain or not we are that one factor causes, or suppresses, another. Values of factors produced in analysis represent how strongly caused (or ‘activated’) they are by changes in other factors based on our level of uncertainty about whether causal links in a map exist or not. Essentially, how certain we are that changes in some factors would cause changes in others, not how large those changes might be. The output of analysing a causal FCM is a list of numerical values of factors under different scenarios with each number representing how relatively strongly caused we believe a factor to be. It aims to answer the question: if we change something in the system, what implications does it have for whether the other factors in the system change?

The second approach, which we will refer to as ‘dynamical’ (again, after Helfgott et al., 2015), considers the propagation of effects of one factor

on another, producing a simple dynamical representation of the relative magnitude of changes in factor values, that is, a plot of all factor values at each iteration. Here, magnitude of factor value does tell us how relatively large an effect is. It is used to help understand which factors and connections are most important or influenced in a system, with larger factor values interpreted as being more important or influenced/influential, and how changing map structure changes that. It can also be used to gauge how relative changes in factor values might play out dynamically, for example, whether change accelerates, stabilises, or dies away.

The differences between these two approaches to analysis are important but tricky to understand. We spend a lot of time in this section unpicking them and return to them in the history of FCM section. Table 6.1 summarises the key differences.

In both the ‘dynamical’ and ‘causal’ approaches, the factors and edges are assigned numerical values. These values and their meanings are different depending on which analysis approach is being used. If the dynamical approach is being used, factor values can take any real value. Depending on the calculations used, the values of factors can represent the size of a change in the factor values or the actual magnitude of those concepts (their initial value plus the change). Commonly, change is explored, and most factor values are initially set as zero, representing a baseline starting position (more on this in a moment). For edges in a dynamical approach, the value will usually be between -1 and $+1$, often in categories corresponding to ‘weak’, ‘medium’, or ‘strong’, which normally represents the strength of the effect of the causal influence (i.e. a value of $+0.5$ would mean if the source of the arrow has a value of 1 , the target of the arrow will go up by 0.5 units, if the value were -0.3 , with the source factor at 1 , the target would go down by 0.3 , and so on). (Although edge values could be given any real value without causing problems in calculation.)

In the causal approach, edge values are constrained to be between -1 and 1 , while factor values are constrained to be between 0 and 1 , or sometimes -1 and 1 . A factor with a value of 1 is fully activated or caused, a factor with value 0 is not activated or caused. Values in between represent how certain we are that a factor is being caused by its inputs. Formally, a value of 0.5 is the most ambiguous, as it means that we have no information about whether a factor is being caused or not. The value of edges represents our certainty about whether factor A causes factor B. If the magnitude of the value on the link is 1 , we are certain that A causes B; however, the lower the magnitude, the less certain we are that A actually causes B. The sign on the link represents positive or negative causation: a

Table 6.1 Comparing the ‘causal’ and ‘dynamical’ approaches to analysing FCMs

<i>Issue</i>	<i>Causal approach</i>	<i>Dynamical approach</i>
<i>Interpretation of arrows</i>	Represent our certainty about whether factor A causes or suppresses change in factor B	Represent the magnitude of influence of factor A on factor B
<i>Initial values of factors/concepts</i>	Set to either 0 or 1 depending on role in analysis/scenario	Can take any real value, but often set to 0, or to 1 if a driver or investigating impact of change in that factor
<i>Value of edges</i>	Must take a value between -1 and 1	Can take any real value, but typically kept between -1 and 1
<i>Intuition behind the purpose of analysis</i>	Assessing the strength of certainty that factors are caused or suppressed by changes in the system. Considers, if we change something in the system, what implications does it have for the causation of other factors in the system	Assessing the relative changes in the magnitude of factor values in the system under different change scenarios. Considers how much (relatively) will magnitudes of different factors be affected by change and thus how much they are influenced or influence others
<i>Key constraints in analysis</i>	Thresholding or squashing function must be applied to factor values to keep them between 0 and 1 or -1 and 1	Thresholding function not required
<i>Outputs of analysis</i>	Ranking of factors in terms of how much they are actively caused or suppressed by a change in the system	Plot showing the relative values of factors through ‘iterations’ of the analysis. Or ranking of factors in terms of how much they influence or are influenced
<i>Pros</i>	Consistent with original approach	Intuitive and appealing output. More intuitive map building
<i>Cons</i>	Counter-intuitive interpretation of arrows and factor values. Does not tell us anything about how much something might change. Results sensitive to form of squashing function chosen	In practice, has created confusion about the appropriateness of interpreting the edges and values in this way. Dynamical output can be misinterpreted as a simulation
<i>Background</i>	Reflects original FCM maths and interpretation in Kosko (1986)	Reflects wide adoption of FCM in participatory mode with a more intuitive interpretation of maps

Source: Authors’ creation

positive link from A to B means that we believe A causes B to happen, and a negative sign means that A causes B not to happen.

Figure 6.1 shows an example of an FCM representing deforestation in the Brazilian Amazon. This map is presented in Kok (2009) and is developed to explore possible futures of deforestation in the Amazon. For comparison, Fig. 6.2 shows an FCM of the UK Humber region bio-based economy from some of our work (Penn et al., 2013). Both maps were constructed to fit with a dynamical approach. We can see how the maps look quite different but contain much the same sort of information. The Humber map is labelled with connections as positive or negative, and then weak, medium, or strong. The Amazon map has the numerical values annotated next to the arrows. Note that drivers (i.e. factors with no incoming connections from within the map) are here given a self-reinforcing loop to prevent their value going to zero during analysis. A map developed for a causal approach would look similar, but without self-reinforcing loops.

Analysis of maps usually involves comparing a baseline scenario, that is, the map as it is under the influence of external drivers or just its own

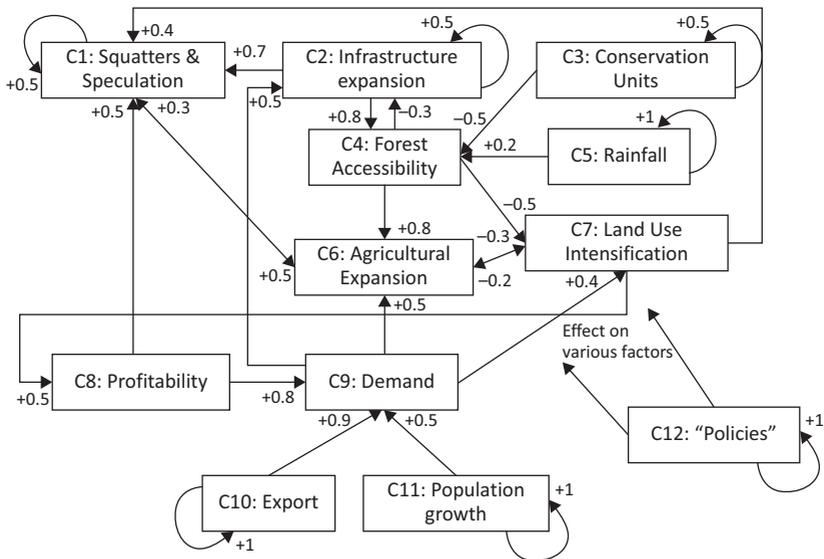


Fig. 6.1 Fuzzy Cognitive Map of deforestation in the Brazilian Amazon. (Source: Kok (2009))

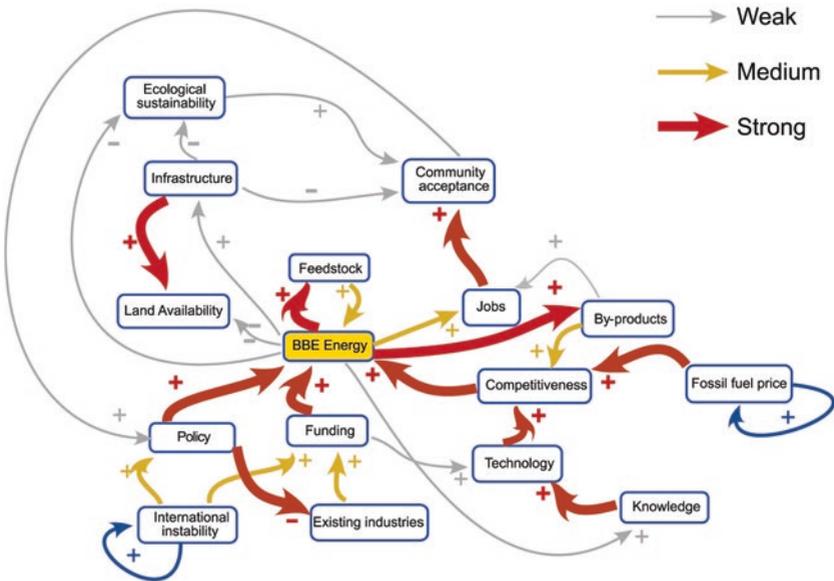


Fig. 6.2 Fuzzy Cognitive Map of the UK Humber region bio-based economy. (Source: Penn et al. (2013))

structure, with additional scenarios corresponding to changes, such as new external pressures or interventions. To do this, the map is transformed into a matrix corresponding to the links between factors and the edge weights (i.e. a table with the factors as both rows and columns, and the values on the edges that connect them in the cells where they ‘cross’). This is the ‘weight matrix’. The factors are all given a starting value, and by repeatedly multiplying and updating this list of factor values by the weight matrix, we get an output of the magnitude of the total influences of all the factors feeding directly into any other factor. Don’t worry if this sounds complicated: it is; the FCM software discussed below will do this for you. Importantly, in the causal approach, output factor values are often modified with a thresholding or squashing function to keep values between 0 and 1 or -1 and 1. This is not required in the dynamical approach.

For the most part, rather than exploring how the actual values of factors or concepts change, FCM analysis explores how changes in a few factors propagate through the system. In the dynamical approach, a baseline scenario is often then produced by setting the initial values of any drivers in

the map to 1 or less and the other values to zero. The model output is then calculated, propagating the change through the map using the values of edges and factors. This process of updating factors' values based on other factors' value changes and edge values is iterated until such point that a stable pattern can be seen in the values of factors. That is either values no longer change, or they change in a repeating cycle.

An example of the types of output this process produces can be seen in Fig. 6.3. The top left plot is the output from the map shown in Fig. 6.1; the other three plots are outputs from slightly modified versions of this map structure intended to represent different policy scenarios. This output is dynamic in the sense that it is based on propagating changes to factor values; however, it is important to remember it does not represent changes through time. Rather, FCM practitioners often refer to changes through 'iterations' of the map. As this is a 'semi-quantitative' method, interpretation involves comparing only the relative ranking of factor values.

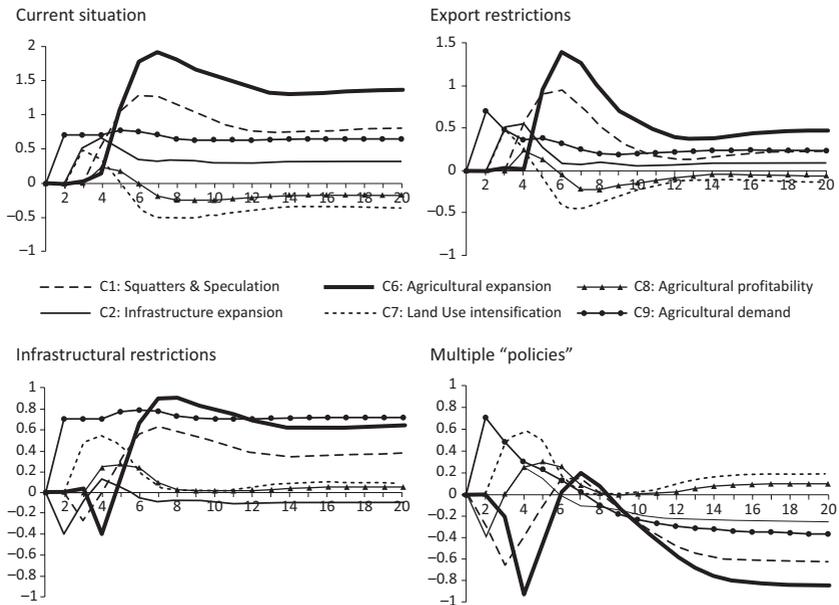


Fig. 6.3 Example outputs from an FCM dynamic analysis. (Source Kok (2009))

There tend to be four types of outputs of the dynamic analysis: (i) the factor values all change and stabilise at new values; (ii) the factor values change and continue in one direction in a ‘runaway change’-type scenario; (iii) values change but return to zero; or (iv) values change up and down in cycles. Looking for these types or calibrating the model to produce them can be a useful discussion tool in workshops. Calibration to achieve this involves tweaking edge weights, in particular, modifying the strength of short feedback loops which have a strong impact on the dynamics. Sensitivity testing may be done to determine which links cause the system to destabilise, but often an experienced modeller will have a sense of what sort of structure will stabilise or not and keep this in mind during map construction. This stage is the most technically demanding, and we recommend you practise and ensure familiarity with it before building maps you plan to use in your research.

In the causal approach, analysis tends to be focused on comparing the impacts of changes in the map on relative factor values. The factors are initialised, and the model run to equilibrium to produce a baseline. To examine different scenarios, one or more of the factors in the map has its initial value increased or decreased and ‘clamped’ at this new value (i.e. held artificially at this point rather than changing dynamically). This represents an external change or an intervention. An output is produced of the relative change in final factor values under a given scenario, compared to the baseline, and all constrained by the squashing function. Again, the structure of the map is often modified to add in an intervention or change and its connections to the extant system. This essentially shows how change propagates through the system as perceived by stakeholders but gives no idea of possible system dynamics.

Combined, the static map of factors, edges, and their values, and the dynamic or causal analysis based on these values is what FCM is. Terminology is fairly settled, there are no other names for the method which are regularly used, and though there can be differences in the names of concepts/factors/variables, and edges/connections, these are normally obvious and make no difference to the use of the method.

The actual practice of FCM does have some more important variety which can go unreported or underappreciated. The analysis can be done in different ways, as we have described above. Moving to building maps, this can be done in one of three ways. Researchers can build them themselves on whatever evidence or input they deem appropriate; multiple maps can be built in individual interviews with stakeholders, and then

combined; or one map can be built in a workshop setting with multiple stakeholders. These modes can be combined in different ways, for example, with a researcher starting a map, then running a workshop, and then conducting some follow-up interviews.

During the construction of a map there are also subtle but important differences on how a map is started and constructed, which can have profound impacts on the nature of maps built. For example, we might have a list of twenty factors, and brainstorm in a workshop the most important connections ad hoc (i.e. asking stakeholders to start where they want), drawing and building out as we connect factors; or we may systematically go through each pair of factors in a table format and consider if there is a connection. Both approaches are valid, have pros and cons, but will result in different maps. More experienced FCM facilitators will also (sometimes subconsciously) guide the process in ways which avoid potential pitfalls later (more on this below).

HOW DO YOU DO FUZZY COGNITIVE MAPPING?

There are three basic stages to building FCM maps. First, we develop a list of the factors to go in the map; second, we construct the map and its connections; and third, we produce the analysis and interpret it with stakeholders or users. Let's consider each of these in turn:

1. **Develop list of factors to include:** This needs to be a well-thought through list of factors which are expressed as variables. Before the list can be made, a decision and agreement must be reached on what the map will be of, that is, what is the system. This system definition can be very difficult. Though factors can be anything, they normally need to have some level of comparability in their abstraction or simplification. In practice, it is useful to limit the number of factors at this stage to twenty, simply to reduce the time required to specify connections at the next stage. In multiple individual interviews, we might provide a list we create ourselves, or let interviewees build their own map with their own factors and then synthesise the maps and factors individuals created, decide when factors are the same and pick a name, remove duplicates, or work out what to do where there are differences. This process is like coding qualitative interviews where we look for themes; it requires a lot of researcher judgement. Guiding interviews based on what has happened in past interviews

can reduce the potential for large differences and difficulty here. In workshops, this stage is normally the most time-consuming as participants must brainstorm, explain their thinking, then agree on choosing, merging, and removing duplicates from a final list. It is common to spend a large proportion of your workshop time on this stage.

2. **Construct the map:** now we can specify the connections between factors, including where arrows should go, discussion of what the process is behind them (sometimes this can uncover confusion, revealing the need for factor redefinition or the need for other arrows or factors), the direction of the arrow, whether it is positive or negative (i.e. increasing or decreasing), and what its strength is (i.e. the value of the edge weight). We do not collect the actual value or strength of the connection until the end when all connections have been drawn, so we can make sure they are all decided relative to each other, often by asking participants to rank them by strength. There can be a lot of confusion at this stage around the meaning of positive and negative connections. The method uses these terms in a mathematical sense (i.e. positive means they move together, as one factor goes up, so does the other, or as one goes down, so does the other; and negative means they move inversely, as one goes up, the other goes down, or as one goes down the other goes up), but stakeholders often use them in a normative sense (i.e. a positive influence is a ‘good thing’, or a negative influence is a ‘bad thing’). It is also perfectly normal at this stage to relabel, add, or remove factors, as stakeholders’ thinking develops.
3. **Conducting and interpreting analysis:** now that we have the static map with factor values and edge values, we can perform analysis using either of the causal or dynamical approaches described above. Because of the importance of the distinction in these two types of analysis, we have described in detail how they are done above. Here, we focus on how these are used in a participatory mode with stakeholders. Analysis in this approach, should be used as the starting point of a discussion rather than a prediction of what state the system will move to. Stakeholders can be asked whether the analysis corresponds to their understanding of the causal effects of the system as they have described it. Does this make sense to them? Do they have intuition about why is it happening? The outputs of FCM analysis in some sense provide a summary or extrapolation of the

causal thinking that the map contains. Thinking through these results considering how their own mental models of causal structure are producing them allows stakeholders to reflect on, question, and confirm or change their beliefs about this structure.

To start doing FCM you need to make a choice about the materials you use to build the map, and then the software you use to visualise and analyse it. At the building stage, the main decision is whether to use pen and paper (and probably post-it notes), or to use software straight away. Using software can seem like an efficiency saving for a map which you will need to digitise at some point; however, it comes at a big potential cost of engagement and inclusiveness when building maps in a group. Using software excludes people who are not confident using computers or unfamiliar software, and if the facilitator operates the software alone, this makes them a bottleneck on the process. In our experience, participants will likely have lower levels of engagement and discussion building a map on a computer, but there will be exceptions to this.

Once you come to digitising the map there are a variety of software options including:

- **General purpose diagramming software:** For example, [diagrams.net](#), Visio, or yEd. There are a huge number of options here. All will be able to build the boxes-and-edges structure (hopefully, in an aesthetically pleasing way), but few, if any, will be able to run the analysis. Nonetheless, it can be useful to have a high-quality drawing of the static version of the map.
- **Spreadsheet software:** It is possible to use spreadsheet software (e.g. MS Excel, LibreOffice Calc) to implement the dynamic analysis of a map with the formula functionality. This will also quickly produce new plots for you, once you have set them up. However, you won't get a nice visualisation of the map itself.
- **FCMapper:** This is a purpose-built piece of software developed by researchers using FCM. It is available from the FCMappers website ([fcmappers.net](#)).
- **Mental Modeller:** This is a purpose-built free browser-based software for FCM and has been used for work in many academic publications. It is based on the causal approach described above.
- **R packages, FCMapper and FCM:** For those who use R, there are a couple of packages to support developing FCM maps. There is a

steep learning curve to using R, so we only recommend this to people who already use R, or who have the time to spend on learning it and are certain they want to use FCM. Given R's flexibility, power, and large user base, it is likely to offer the most functionality for building FCM maps, though we have not used it for this task ourselves.

COMMON ISSUES AND 'TRICKS OF THE TRADE'

There are innumerable issues and questions you can come up against in an FCM project, too many to cover here in detail. Instead, we will consider some of the more salient and conceptual issues that apply to FCM.

With FCM it is often the dynamic analysis that users and stakeholders are most drawn to. Although all FCM practitioners refer to their method as 'semi-quantitative' and emphasise that it is not a simulation approach, in practice, there is always a temptation to over-interpret the model output. The analysis can seem like magic (especially when done live in a workshop) and offer a false sense of certainty, truth, and scientific rigour. The outputs of any dynamic model are sensitive to the assumptions in that model, and the dynamic analysis of an FCM map built based on a participatory process is the same, but those assumptions have come from a place of (quick and dirty) consensus building and group deliberation, rather than cold methodical modelling. Additionally, the simple nature of the mathematics itself, whilst allowing rapid modelling, can hide unexpected pitfalls. In a dynamical FCM, edge weights and the presence and strength of feedback loops drive the output. The process of model calibration, altering these connections to obtain an equilibrium is intended to make the model output more interpretable by stakeholders, not to build a more accurate simulation.

If, on the other hand, a causal FCM approach is being used, with threshold functions on the factor values, then the model output is extremely sensitive to the mathematical form of these functions. In fact, when using this mathematical formulation, changing the form of the threshold function can have more of an impact on the results than the map structure, even to the degree of reversing model results. It is crucial therefore to do extensive sensitivity analysis when using a causal approach. We do not mean to suggest the model output has no value; it does. However, it is vital that it is used in the right way (i.e. as a discussion and thinking

tool rather than as a forecast of what might happen), and those being shown it understand these caveats properly.

These, and other, common issues can normally be dealt with careful planning, an eye for detail, and a reflexive approach. Iterating through map building, analysis, and interpretation will also be invaluable in ironing out issues. Iteration generally is a key theme to many of the approaches in this book—you will be sick of us talking about it soon!

There are also some tricks of the trade for FCM which can unlock potential issues. The first, is to think about the analysis right from the start of map building. This is tricky at the beginning, but with practice you start to intuit whether a map structure will produce interesting outputs or not. As you build, look for feedbacks; are there any? Are there many short loops (i.e. with factor A affecting B and B affecting A)? Ideally you want to have a map with feedback loops, but these need to be questioned carefully so they reflect reality and beliefs as best as possible, not just dropped into the map without much thought. Direct loops should always be questioned, are there intermediate factors which we should route a feedback via? A short, direct feedback will likely have a strong effect on the dynamic analysis, can we specify it so there are more steps?

The ability of FCM maps to include different types of factors is appealing. However, this can lead to issues if factors operate on different timescales. For example, if factor A influences factor B over many years, and factor C also influences B, but on a daily basis, the analysis of the map will treat them in the same way, which may produce misleading outputs. One solution may be to make the long-term influence take a small value though this won't always make sense. A more reliable solution is to try to redefine the boundary of the map, and thus the system being considered, so that factors with very different timescales do not have to be artificially included together.

Much of the value of FCM is in a process of capturing and then challenging stakeholders' mental models. In the dynamical approach, often altering the map live in a workshop to expose ways in which simple changes in system structure can alter outputs. Although this takes much experience to do live, a simple trick to try is to check some of the key outcome factors in the map and to ensure these are not represented as sinks (i.e. only with incoming arrows). Producing maps with this structure is common but leads to boring dynamics and is unlikely to reflect reality.

WHAT IS FUZZY COGNITIVE MAPPING GOOD AND BAD AT?

FCM's biggest strength (but also arguably the biggest risk in using it) is in its intuitive and quick descriptions of a system's structure and semi-quantitative outputs. It does not require data or empirical evidence to inform or validate the map and its analysis. It is a form of quantitative storytelling, bringing together stakeholders and their narratives in a way which allows stakeholders to question and perhaps change their own assumptions about the system. The quick and dirty approach makes it ideal for a workshop context and use with stakeholders, and its ability to capture qualitative and quantitative aspects of a system mean that anything that matters on the ground can be included, making the method inclusive. The dynamical analysis, as well as being engaging and exciting to see, extrapolates and makes visible to stakeholders the logical implications of their beliefs about system structure, as it shows what factors are more highly driven, influenced, or caused by external changes or interventions. Essentially allowing them to ask, 'does this factor or change actually cause what I think it does according to what I believe the system structure to be?'. This can lead to profound learning experiences.

The nature of the analysis required to work in this quick and dirty way, however, is the source of many problems. It creates the possibility of the analysis being easily misapplied or misunderstood. At worst, meaningless, but quantitative, analysis could be taken as truth by users and stakeholders. Although the dynamical analysis is often cited as a key hook, many practitioners say that the most important output is the map itself.

A BRIEF HISTORY OF FUZZY COGNITIVE MAPPING

FCM first appeared with Bart Kosko's (1986) paper outlining a modified approach to cognitive maps (themselves first outlined by Axelrod, 1976), applying fuzzy causal functions to connections (i.e. using -1 to 1 to represent the certainty of presence of a + or - causal link). Cognitive maps had been around for a little while and had been applied in a variety of domains; they had binary values for connections (i.e. either on or off, there or not) and were among the first representations of causal connections between factors identified by stakeholders, rather than researchers alone. Kosko's contribution in adding the 'fuzzy' was two-fold; to allow a more nuanced representation of causal reality, and to allow the maps to be 'computed'

(i.e. with numerical values, to allow values of factors and edges to be combined to calculate the outcome of the values and the map structure).

Initially, most FCM research was reported in technical journals, with focus on the technical and mathematical details and methodological development. This meant FCM was slow to reach new audiences and potential application areas. While there is still much highly technical FCM work, since 2000, there has been both an increase in the number of published FCM works, and a spread into domain journals and a variety of fields, from social sciences through medicine and natural sciences. There does appear to be a clear divide between (i) work which uses FCM in a participatory mode, treating maps primarily as discussion and learning tools; and (ii) work which focuses on model discovery/inference/learning from data and treats maps as a form of neural network. Jetter and Kok (2014) provide a detailed history, with reference to many examples, for readers wanting to explore this history more fully.

Over time, however, as this method has been used more in participatory contexts, an ambiguity about whether certainty of cause and activation, or strength of causal link and magnitude of change, in factors is being investigated. Many authors use a causal mathematical modelling approach, but discuss strength of connections and size of effects, or switch between causal and dynamical terminology. This is potentially problematic as we are not actually calculating how values of factors change with this formulation and so a danger exists that a non-sensical model is produced and hence misinterpreted. A consistent causal approach could be used within stakeholder workshops by mapping certainty regarding cause rather than strength of effect on links. However, the ideas of factors being caused or not caused are by no means as intuitive as thinking about factors having an actual value or link weights representing a size of effect of one factor on another. This could cause issues within workshops when soliciting values for the map and indeed seems to cause issues in the literature. These issues are discussed in detail and clearly explained in the excellent working paper by Helfgott et al. (2015).

GETTING STARTED WITH FUZZY COGNITIVE MAPPING YOURSELF

There is a noticeable dearth of resources and guides for getting started with FCM. However, we would recommend the following reading:

- Jetter and Kok (2014) provide a relatively accessible history and introduction to FCM, as well as a detailed discussion of how to ‘do it’, including consideration of how to design a process, not just the method itself.
- Özesmi and Özesmi (2004) provide a detailed description (which sometimes errs on the technical side) of how to apply their (highly cited) approach to using FCM, including exploring policy scenarios.
- <https://www.mentalmodeller.com/> as well as access to the software, the mental modeller website introduces FCM and several case studies of its use.
- Felix et al. (2019) provide a more technical review of FCM, including discussion of software options.

Beyond getting a firmer and more detailed grip on how to use the method, we would recommend two things. First, decide on the materials and software you are going to use to help you do FCM. This is an essential choice and making it will allow you to experiment with the method (either making ‘test’ maps in small groups with colleagues, or playing with different software) to get a real hands-on feel for it. Second, you need to find the windows of opportunity to apply FCM usefully. If you want to use it with stakeholders, speak to them and find out how it might be of value, what are the questions they have that might be amenable to FCM, what are the processes and workflows which an FCM project might feed into.

REFERENCES

- Axelrod, R. (Ed.). (1976). *Structure of decision: The cognitive maps of political elites*. Princeton University Press. <https://doi.org/10.2307/2616237>
- Felix, G., Nápoles, G., Falcon, R., Froelich, W., Vanhoof, K., & Bello, R. (2019). A review on methods and software for Fuzzy Cognitive Maps. *Artificial Intelligence Review*, 52(3), 1707–1737. <https://doi.org/10.1007/s10462-017-9575-1>
- Helfgott, A., Lord, S., Bean, N., Wildenberg, M., Gray, S., Gray, S., Vervoort, J., Kok, K., & Ingram, J. (2015). Clarifying fuzziness: Fuzzy Cognitive Maps, neural networks and system dynamics models in participatory social and environmental decision-aiding processes. *Processes EUIP7 TRANSMANGO Working Paper 1*.
- Jetter, A. J., & Kok, K. (2014). Fuzzy Cognitive Maps for futures studies—A methodological assessment of concepts and methods. *Futures*, 61, 45–57. <https://doi.org/10.1016/j.futures.2014.05.002>

- Kok, K. (2009). The potential of Fuzzy Cognitive Maps for semi-quantitative scenario development, with an example from Brazil. *Global Environmental Change*, 19(1), 122–133. <https://doi.org/10.1016/j.gloenvcha.2008.08.003>
- Kosko, B. (1986). Fuzzy Cognitive Maps. *International Journal of Man-Machine Studies*, 24(1), 65–75. [https://doi.org/10.1016/S0020-7373\(86\)80040-2](https://doi.org/10.1016/S0020-7373(86)80040-2)
- Özesmi, U., & Özesmi, S. L. (2004). Ecological models based on people's knowledge: A multi-step Fuzzy Cognitive Mapping approach. *Ecological Modelling*, 176(1–2), 43–64. <https://doi.org/10.1016/j.ecolmodel.2003.10.027>
- Penn, A. S., Knight, C. J. K., Lloyd, D. J. B., Avitabile, D., Kok, K., Schiller, F., Woodward, A., Druckman, A., & Basson, L. (2013). Participatory development and analysis of a Fuzzy Cognitive Map of the establishment of a bio-based economy in the Humber region. *PLoS ONE*, 8(11), e78319. <https://doi.org/10.1371/journal.pone.0078319>

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

