Sara Lumbreras and Lluis Oviedo

Belief networks as complex systems

ENGLISH

ABSTRACT

There has been extensive work on understanding belief, from a psychological, philosophical and neurobiological perspective. Meanwhile, artificial intelligence has produced compelling developments that can enrich and update the brain-as-a-computer metaphor and has tried to better represent beliefs as cognitive probabilistic processes. In parallel, there has been a surge of research in Complexity Sciences, with applications ranging from Medicine to Finance. Some authors have already linked the connected nature of belief to the behaviour of complex networks. We would like to expand this approach to understand belief as a complex system with the main functions of providing a model of the world – including the individual and her surroundings – and producing guidelines for action. The complex-system perspective allows us to understand some of the properties of belief systems in a comprehensive manner, which many authors have begun to study in isolation. Notably, this provides a framework to study the important phenomena of belief formation and change as processes of emergence and adaptation. In this exploratory paper, we propose an outline for this framework for this study.

DEUTSCH

Glaube wurde bereits vielfach aus psychologischer, philosophischer und neurobiologischer Perspektive untersucht. Inzwischen gab es im Bereich der künstlichen Intelligenz beeindruckende Fortschritte, die die metaphorische Bezeichnung des Gehirns als Computer aktuell und relevant erscheinen lassen und Versuche nahelegen, Glaube als kognitiven probabilistischen Prozess darzustellen. Gleichzeitig wurde die Komplexitätsforschung intensiviert, deren Anwendungen von der Medizin bis zum Finanzbereich reichen. Einige AutorInnen stellten bereits Zusammenhänge zwischen der Natur des Glaubens und dem Ver-

halten komplexer Netzwerke her. Der Ansatz, Glaube als ein komplexes System zu verstehen, dessen wichtigste Funktionen die Bereitstellung eines Weltmodells – einschließlich des Individuums und seiner Umwelt – und die Erstellung von Handlungsrichtlinien darstellen, soll hier weitergeführt werden. Ein solcher systemtheoretischer Zugang ermöglicht es, einige in der Forschung bislang zumeist isoliert betrachtete Eigenschaften von Glaubenssystemen ganzheitlicher zu verstehen. Insbesondere entwirft dieser Zugang einen Rahmen für die Untersuchung so wichtiger Phänomene wie der Glaubensbildung und der Glaubensänderung als emergente und adaptive Prozesse. Der vorliegende Artikel skizziert den Entwurf eines solchen Rahmens.

BIOGRAPHIES

Sara Lumbreras is a professor at the ICAI School of Engineering (Madrid, Spain). Her research focuses on the development and application of decision support techniques for complex problems, mainly applied to energy, finance and health care. She specializes in mathematical optimization and machine learning. In addition, she is a member of the board of the Chair of Science, Technology and Religion at Universidad Pontificia Comillas, where she studies the evolving links between technology and society.
E-Mail: s.lumbreras@comillas.edu

Lluis Oviedo, born in Spain, is full professor of Theological Anthropology at Antonianum University, Rome; and invited professor in the Theological Institute of Murcia (Spain) for questions of religion, society and culture. His research focuses on new scientific study of religion and its theological impact, Christian and scientific anthropologies, and issues of secularization and religious social dynamics.
E-Mail: loviedo@antonianum.eu

KEY WORDS

Artificial Intelligence; Artificial Neural Networks (ANN); belief; belief adaptation; belief networks; complex systems; decision making; emergence; pattern recognition; Reinforcement Learning (RL)

1 Introduction: Credition as a fundamental brain function

Recently, credition – the process of believing – has been defined as the functions that enable somebody to trust her inner probabilistic representations. Credition includes perception and valuation, and guides action by means of reciprocating feedback involving learning. We understand credition as a basic cognitive function, essential to understand the human mind and human behaviour. As such, it is receiving increased interest in the literature (Angel et al. 2017). We cannot proceed without beliefs.

Credition is essential to understand the human mind and human behaviour.

From this general framework, the task that the study of the believing process faces is to find models that allow for a better representation of this rather enigmatic process. There have been several attempts to model the believing process over the last years. The credition model emerges as a general platform that allows for many applications. Recent developments in Artificial Intelligence (AI), the study of probabilities and complex systems, might offer clues to further the understanding of such a process and to build more accurate models that are able to capture its complexity. In any case, our exploration will show both, the strengths and convergence points between AI systems and belief systems, and the diverging points that arise in that contrast.

This exploratory paper proposes an outline for a framework in which we can understand belief and the process of credition, together with their properties, from a complex-systems perspective.

2 The old brain-computer metaphor is dead. Long live the new one: incorporating the recent developments of Artificial Intelligence

The traditional interpretation of the brain-as-a-computer has not stood the test of time and is seen by most specialists as outdated. The animal brain is much more than a deterministic device following some defined code. Even if we accept that the brain's function is to maximize the survival and reproduction possibilities of its owner, this is nonetheless often accomplished as an act of pure creativity. The brain builds a model of the world that constantly integrates new evidence and provides an explanation of ongoing experiences to guide future actions. All these functions are supported by beliefs and the process of credition.

The recent successes of AI have injected renewed strength into the computer metaphor. While the old computer metaphor, which attempted to assimilate the mind to a particularly efficient calculator, has proven unsuccessful, alternative approaches to this simile are starting to emerge, based on connectionist models rather than serial ones and flexibility and adaptability as opposed to predictability and design. Many are the thinkers who see in the recent achievements of AI the basis to substantiate claims for the possibility of creating even artificial consciousness (Kurzweil 2012).

The remainder of this section explores the two most prominent AI techniques in this context: Artificial Neural Networks (ANNs) and Reinforcement Learning (RL).

Artificial Neural Networks and Reinforcement Learning

Since many newly published papers in the field of AI studies deal with 'beliefs', it appears that such an approach might become fruitful when we try to untangle the intricacies of the believing process. Just as an example, browsing for the term 'belief' in the titles of research articles in the Journal *Artificial Intelligence*, we get 739 entries, i. e. published titles that include the word 'belief'. Among published titles in journals or edited books in the field of AI, we find for example: Belief and truth in hypothesised behaviours; Group Decision Making via Probabilistic Belief Merging; Lifted first-order belief propagation; Probabilistic Belief Embedding for Knowledge Base Completion; Causal Basis for Probabilistic Belief Change: Distance vs. Closeness; Probabilistic Belief Revision via Similarity of Worlds Modulo Evidence; Belief Systems and Partial Spaces.

The brain as a pattern-recognition system

Pattern recognition is arguably one of the main functions of the brain, and a particular field where the success of AI has been incontestable. Artificial Neural Networks are computer systems vaguely inspired by the neural processes in animal brains. We recommend the introductory work by Hassoun (1995) for a complete working guide, but will proceed to provide a stylized description of this tool for the purposes of supporting our framework.

Networks are composed of several units (*artificial neurons*) that work together to perform tasks defined by the programmer, such as classification or prediction. The programmer does not input any specific rules into the network. On the contrary, the system "learns" by means of example. For instance, an ANN can be trained to identify pictures of cars, without knowing that cars have four wheels, seats, or a trunk. Instead, the network receives pictures that have been labelled as "car" or "not a car", and infers what the underlying characteristics of automobiles are.

The programmer does not input specific rules into the network.
The system "learns" by means of example.

This is what we call *supervised learning*, as the system receives items that have been correctly classified. In other words: the network recognizes the pattern that appears in the examples it receives. ANNs have been remarkably successful in performing difficult tasks such as computer vision, speed recognition, machine translation and medical diagnosis. All these applications correspond to classification problems, where we need to identify which set of categories a new observation belongs to. There is a second, no less important type of problem known as forecasting. In forecasting, the network detects the patterns that underlie the time-dependent evolution of a variable and predict their unfolding. ANNs have also excelled at this task.

There are many different flavours of networks, which have been proven to have varying strengths. Normally, these networks are structured in several layers, where some of the units are in direct contact with the input they receive, others constitute the output and the remaining ones stay hidden. Each of the neurons receives an input from the set of neurones that are connected to it and it uses this input to generate a single output by means of a relatively simple function, usually just a linear combination of the inputs and some weights. This linear combination is passed through an *activation* function that maps it into the interval [0,1]. The hyperbolic tangent or the sigmoidal function are some of the main activation functions used in ANNs. These simple calculations provide the framework for the ANN. The weights that will define each neuron are calculated by the application of what is known as a *training algorithm*. All training algorithms start by allocating random starting weights to the network that will be progressively adjusted taking into account the errors that they create in the outcome. *Backpropagation* is the method that calculates how a current error is related to each of the weights and how they should be adjusted. It is efficient and can be executed with short computation times, so that it is possible to quickly obtain weights that lead to very few errors. This would mean, in our example, that when the ANN is shown new pictures, it can accurately distinguish between car and non-car elements.

In general, when there are more hidden layers in an ANN, the network can carry out more complicated processes. *Convolutional* ANNs, and other types of *deep networks* (in the sense of having many layers) have proven to be especially efficient in performing complicated tasks such as image processing. Here, the layers can be understood as providing a hierarchical structure for the patterns that are recognized abstraction. This hierarchical structure has been linked to the functioning of the brain, which some authors (notably Ray Kurzweil) have defined as a pattern-recognition system (Kurzweil 2012). According to Kurzweil, the brain is structured in pattern-recognition units that are activated when exposed to a similar stimulus. For instance, they might recognize a small black line on a white paper. Then, several of these patterns combined might be recognized as a larger pattern. Following this example, they might respond to the letter "A". Several of these letters might form the word "APPLE", which triggers the concept of the fruit, and so on.

We can understand a different problem, forecasting, as a form of pattern recognition. If we feed time series of rain inflows to an ANN, it can learn to predict how much rain will come next, because it recognizes the patterns in the data that reflect the season or any particularities of the region that is being analysed.

There are functions of the brain that could be well described as pattern recognition.

The renewed brain-computer metaphor assimilates the brain to a pattern-recognition system. There is merit in this metaphor, which appears closer to the truth than the dated brain-computer simile. The animal brain is far from following a deterministic code. There are indeed some functions of the brain that could be well described as pattern recognition. Belief systems can also have a pattern recognition function and act as classifiers, a task that can be understood as establishing what something is and what it is not. In addition, belief systems can also be used for forecasting, as they provide a model for the world that help us anticipate our actions.

Reinforcement Learning as a metaphor for decision making

Another AI technique that can enrich the brain-computer metaphor is Reinforcement Learning (RL). RL simulates an agent that can take actions in her environment. These actions can lead to a reward. The agent learns, by trial and error, the consequences of her actions. Then, she can take the decision that is optimal for her situation (we will refer to this as her state).[1]

Reinforcement Learning provides a mechanism to understand how, from meta-goals, we can derive intermediate goals and optimal strategies.

We define "optimal" according to a given objective, described by means of an objective mathematical formula. The power of RL and *Dynamic Control*, its underlying technique, is that by defining only the value of the different final outcomes, it can derive the value of the intermediate outcomes, so that the method arrives at the optimal strategy at each stage even if it is not the final one, as the final value cascades into the nearer ones. For instance, if we apply RL to playing chess, the final outcome can be winning or not. The next-to-last states can be described in terms of their probability of resulting in a winning situation. Then, the ones before can be evaluated in terms of the next ones, and so on. Reinforcement Learning therefore provides a basic mechanism to understand how, from meta-goals, we can derive intermediate goals and optimal strategies.

Belief systems have several functions as has been recognized in the literature (Frank 1977). One of these functions is to evaluate a given state in order to select the next most desirable course of action, so that they have a key role in these processes.

3 The limitations of the pattern-recognition simile

The brain is far more complex than an ANN. A single neuron is an entity of a remarkable complexity, which cannot be reduced to a mere simple mathematical operation. We should remember that even unicellular organisms exhibit complex behaviour: amoebas hunt, protozoa build complex shells from minuscule specks of dust. A living cell is a wonderfully complex being, and we should keep in mind that ANNs are only vaguely based on the workings of the brain; they are not modelled as its equivalent.

1 For more context on these techniques, we refer the reader to Sutton/Barto 2018.

More importantly, belief systems have an internal structure that has not yet been reflected in the workings of ANNs, or machine learning in general. The most important difference we would like to highlight is that belief systems are built as a hierarchy of concepts or categories. A concept is a pattern that is comprehended, that has meaning in the sense of understanding the concepts that relate to it in the hierarchy and that are either at its same level, correspond to a partial aspect of it or correspond to a generalization of itself. The relationships between the concept and other concepts in the network can be known partially or in full and can be subject to ambiguity. Powerful concepts have a simple description and can be applied to a large number of instances, while weaker concepts are more convoluted or ambiguous. Some authors have postulated how belief emergence and change maximizes explanatory power or minimizes cognitive dissonance. This parsimonious quality of belief systems has a key property that has not been noted previously, at least to the best of our knowledge: they are robust with respect to generalization.

Belief systems are robust with respect to generalization.

Any machine-learning engineer working with ANNs to solve classification or forecasting problems knows that their main danger is *overfitting* (Hawkins 2004). Overfitting happens when the ANN learns the examples too well and is not able to generalize. For instance, if most of the pictures of cars we show our network correspond to vehicles photographed on a sunny day, the algorithm might have trouble recognizing a car when the surroundings are dark, as it does not recognize that lighting is not a relevant feature when identifying a car. ANNs (and any other machine-learning technique) do not understand what a car is; they are not able to discern when a characteristic is important and when it is not. They only infer this from the examples they are given. Overfitting is an issue that has to do both with the examples we show them and with the training methodology, for instance taking care to select a representative enough set of instances. The fundamental trouble with machine pattern recognition in general (not only ANNs) is that although the pattern might be recognized correctly, it is not comprehended. The pattern of a car, for a computer-recognized pattern, might be a complex combination of light and dark spots on an image. The network does not need to recognize the wheels or the engine to identify a car, so when the situations vary slightly (such as in the lighting example) they might lead to errors.

In belief systems, concepts are at least partly comprehended in their relationships to others, so they excel at generalization. This is remarkably important in an evolutionary sense, given that humans live in a changing environment where no two situations are precisely alike.

Belief systems provide a filter for experience.
Its categories define what features of reality are relevant.

Another interesting property of belief systems that distinguishes them from the architectures of machine learning is that they provide a *filter for experience*. This filter would have a stronger effect the more ambiguous the situation. Many authors have explored the bidirectional relationship between belief and evidence (Fryer Jr/Harms/Jackson 2019).

The programmer defines the examples that will be received by an ANN, for instance as a grid of shades that encodes an image digitally. However, the categories in our belief systems define what features of reality are relevant. All experience gets filtered through the belief system, and its information is used to dynamically update the belief system.

The change can happen in two ways. The first one corresponds to *belief emergence*, when a pattern of relationships between other concepts is found to appear repeatedly or in a significant manner. Belief emergence is one of the key phenomena in this context and has been subject to extensive study (Keil 1991).

Belief emergence can only be linked to machine pattern recognition in a superficial way because, as explained above, animal pattern recognition is based on concepts and is more robust to generalization. When a pattern is sufficiently important, we create a category and usually give it a name. Then, the relationships between the previously held and the new concept are described (similarity, difference, proximity, etc.). Probably, only the concepts that seem to be close in terms of similarity or proximity are automatically (unconsciously) scanned for their relationship to the new concept. It is plausible that, in a first phase, relationships are proposed based on imperfect memory and are subsequently tested for validation in practical experiments. The human brain seems to be especially adapted to detecting subpatterns that help it to generalize. A particularly enlightening example is the Doman method for reading (Doman/Doman 1994). In this method, children are presented with cards that spell out certain words. Their caregiver is supposed to read the whole words to the toddler repeatedly. After enough practice, children not only recognize the words on the cards, but they can also read any other word. They have unconsciously learnt to identify the letters and can generalize in their ability to read but have learnt much more quickly than if the letters and combinations are shown to them directly. This method has been proven to work even with toddlers who cannot yet speak, given that understanding visual information is much easier than articulating words. After all, human beings speak well before understanding grammar, and walk without any conscious idea of the inner workings of their anatomy.

The second type of change in belief systems is *belief change*, where an alternative definition of a concept is found to have more explanatory power or to be less subject to ambiguity (Bendixen 2002). If the new concept definition is evaluated as considerably superior to the one previously held, then the belief changes. We discuss some further details of these processes in connection to complexity theory in the next sections.

An additional missing issue that has not yet been adequately introduced in machine learning is the consideration of logical constraints. However, there have been some very interesting developments in understanding how beliefs can be related to each other through logical constraints and how individuals influence each other (Friedkin et al. 2016). This is especially important: cognitive dissonance emerges when these logical constraints are not respected.

4 The limitations of the Reinforcement Learning perspective

The main issue in the RL perspective is the definition of the objective function, that is, the evaluation of the different outcomes for the agent. When a programmer is developing the system, she creates a mathematical function that expresses her preferences (i. e. what states are preferred compared to others). However, it is not clear how this function would be built in the case of a living organism.

It seems intuitive to choose survival, or reproduction, as the final goal for a living being. Then, some intermediate goals could be chosen by natural selection, given that they will favour survival and reproduction. Evolution may have intuitively selected homeostasis as a process that supports survival. This would include, for example, physical integrity or feeding when hungry.

Then, individuals can derive intermediate goals that get them to a point where keeping homeostasis is easier or more difficult. RL provides an intuitive understanding of how belief systems encode how to derive the impact on the final goals from the intermediate ones. An example of this could be the belonging to a group for many animal species. The group can be critical for survival, so much that this intermediate goal can be ingrained genetically as deeply as the fight for physical integrity. However, the fact that some behaviours are not accepted by the group must, in most cases, be learnt by experience. The individual will learn not to contravene social norms and what those social norms are in her particular context. She will acquire beliefs of what is acceptable and what is not. Let us imagine a human group, where the individual learnt that money (i. e. financial status) is important for belonging to that group. This could cascade to a job being a lower-level goal and respecting the commands she receives as an even lower one.

It should be noted that these processes could be replicated at other levels in a multilevel selection framework. For instance, there is not only selection and learning operating at the level of the individual, but also at the level of kin, group or culture.

Human beings create and morph their objective function in an act of creativity.

However, human beings can indeed contravene social norms, or take decisions that endanger their survival, if they consider that it is the right thing to do. Humans do not only take a defined objective function and learn how to act according to it; they create and morph their objective function in one of the most remarkable acts of creativity. Indeed, values are also defined by means of belief networks.

Most of the interesting features of the renewed brain-computer metaphor can be retained if we study beliefs as a complex system. This perspective can also bring other interesting phenomena to our attention.

5 Beliefs as complex systems

Beliefs cannot be understood individually but only as networks, as has been adequately stated in recent literature. However, a complex-systems approach provides a more comprehensive perspective.

A complex system is an entity composed of many parts that interact with each other, and whose behaviour is difficult to predict although the constituents might be simple to define. Examples of complex systems can be found in biology, sociology or engineering.

A complex system is an entity composed of many parts that interact with each other, and whose behaviour is difficult to predict although the constituents might be simple to define.

Complex systems have distinct properties that arise from their internal relationships, such as nonlinearity, emergence, spontaneous order, adaptation and feedback loops. These properties appear in such a wide variety of fields that they have become the object of independent research, which applies to the system irrespective of its nature.[2] Their power lies in how general these properties are, and how they help explain very different phenomena. For instance, from a complex-system perspective, the behaviour of a flock of birds trying to move ahead and be protected while avoiding collisions could be linked to that of a group of firms trying to develop their companies according to established business models while avoiding excessive competition.

We propose that belief networks can be studied as a complex system, and that this perspective helps to understand their dynamics as it links them to general phenomena that appear in other systems irrespective of their nature.[3] The remainder of this section details some of the most interesting aspects of this approach.

- *Complex systems have goals, and so have belief networks.*
  In particular, belief networks have some defined goals, which have been defined in the literature according to diverging criteria (Frank 1977). We identify at least three defined goals. They filter the world, so that any new evidence is interpreted through the prism of the existing beliefs. In addition, belief systems also provide a model for the world, so that they can explain history (global and personal) and can be used to anticipate the consequences of actions. Finally, they can define what is important, and prioritize, what should or should not be done.

- *Complex systems are open, and so are belief networks.*
  Complex systems must receive input from the exterior world in order to survive. This input can be material or take the form of energy or information. Belief systems correspond to the latter case. They

2 A good introductory text on complexity theory can be found in Mitchell 2009.

3 We follow the outline in Mitchell/Newman (2002) for the properties of complex systems and encourage the reader to refer to it for more detailed information.

are open to interaction with the environment and with others, so that beliefs are adapted in the light of new evidence or through interaction with others (Crocker/Fiske/Taylor 1984; Rodriguez/Bollen/Ahn 2016; Sodian/Zaitchik/Carey 1991).

- *Belief systems are complex.*
  They can be generated from a relatively simple starting set of ideas, but they are not easy to understand or predict. There is a breadth of literature that deals with how to predict belief formation and evolution depending on personal characteristics. Personal features such as analytic cognitive style (Pennycook et al. 2012), feelings of superiority (Toner et al. 2013), or even parenting styles (Ruffman/Perner/Parkin 1999) have been shown to influence belief formation and change.

- *Belief networks, as complex systems, are subject to nonlinear phenomena.*
  Nonlinearity means that the same stimuli do not lead always to the same response. For instance, it takes more information to change beliefs than to confirm them. This bias appears also in many systems under the general name of hysteresis. Hysteresis occurs when the response depends on the history of the system. For example, a magnet may have more than one possible magnetic moment in a given magnetic field, depending on how the field changed in the past. Hysteresis has been studied in physics, chemistry, engineering, biology, and economics, and probably also appears in belief change. In addition, some complex systems can experience very different states with respect to change: a relatively stable state when changes are slow and a "crisis" state when changes can happen rapidly and spread widely. This has been thoroughly studied, again, in magnetism, and has been observed in what is known as the Ising model (Kaneyoshi/Jaščur/Fittipaldi 1993), which describes a magnetic material in terms of its microscopic domains. In belief networks, change is generally difficult but, in times of crisis, the change in one belief can spread to a large number of them. This is true for personal and for social beliefs, where the crisis dynamic could be explained as a paradigm shift (Jones 1977).

- *Emergence:*
  This is probably the key element in the characterization of complex systems. The emergence of new beliefs is a creative phenomenon. New properties emerge from complex systems. This has been highlighted as one of its main properties in Systems Theory. In particular, this creativity should be understood as strong emergence in the sense proposed by Chalmers (2006).

- *Spontaneous order and self-organization:*
  Beliefs organize in more or less consistent and related spheres of influence. They tend to generate in a cohesive manner, with beliefs that have a similar context being closer as well. This means that beliefs about similar things tend to emerge and adapt at relatively close periods of time. In addition, they spontaneously order themselves in a hierarchical structure, which mimics the structure of concepts in our own minds (Kurzweil 2012). Consistent belief systems have been described by some as "attractors", states towards which a system tends to evolve (Goertzel 1995).

- *Adaptation:*
  Belief systems adapt to better fulfil their objectives as described above. The dynamics of belief change have strong dependencies on environmental, personal and social factors (Rodriguez et al. 2016).

6 Concluding remarks

This exploratory paper has reflected on the implications of developments in Artificial Intelligence for the study of beliefs. The reviewed advances provide interesting clues on how beliefs work and how they are structured, but at the same time, they reveal the limits of this model, since belief systems in humans have many specific traits that can be hardly reduced to AI systems. Nevertheless, this approach encourages the development of new models of beliefs as complex systems which follow patterns in AI and probabilities calculation to some extent but cannot be reduced to them.
In all, the ongoing research offers a more in-depth approach to the study of beliefs as specific processes that cannot be easily reduced to other similar cognitive activities, like perception, or evidential cognition, and that play an important role in human life and decisions. The application range of this

approach is very broad, from factual beliefs or simple representations about immediate reality, to values, ideologies and religious beliefs. In all cases, current research proposes a much more cautious approach and to recognize the complexity and cognitive value of such processes, perhaps once neglected or even dismissed as 'secondary', 'derived' or even delusional cognition forms. The study of belief will thus be central to understanding some of the most important processes in human life and decisions.

References

Angel, H.-F. / Oviedo, L. / Paloutzian, R. F. / Runehov, A. L. / Seitz, R. J. (2017), Processes of believing. The acquisition, maintenance, and change in creditions, Cahm: Springer.

Bendixen, L. D. (2002), A process model of epistemic belief change. https://www.researchgate.net/publication/232454198_A_process_model_of_epistemic_belief_change [03.09.2020].

Chalmers, D. J. (2006), Strong and weak emergence, in: Clayton, P. / Davies, P. (eds.), The Re-emergence of Emergence, Oxford: Oxford Univ. Press, 244–256.

Crocker, J. / Fiske, S. T. / Taylor, S. E. (1984), Schematic bases of belief change, in: Eiser, J. R. (ed.), Attitudinal judgment, New York: Springer, 197–226.

Doman, G. / Doman, J. (1994), How to teach your baby to read. The gentle revolution, Avery.

Frank, J. D. (1977), Nature and functions of belief systems. Humanism and transcendental religion, American Psychologist 32, 7, 555.

Friedkin, N. E. / Proskurnikov, A. V. / Tempo, R. / Parsegov, S. E. (2016), Network science on belief system dynamics under logic constraints, Science (New York, N.Y.) 354(6310), 321–326. DOI: 10.1126/science.aag2624.

Fryer Jr, R. G. / Harms, P. / Jackson, M. O. (2019), Updating beliefs when evidence is open to interpretation. Implications for bias and polarization, Journal of the European Economic Association 17, 5, 1470–1501.

Goertzel, B. (1995), Belief systems as attractors, in: Robertson, R. / Combs, A. (eds.), Chaos Theory in Psychology and the Life Sciences, Mahwah, NJ: Erlbaum, 123–134.

Hassoun, M. H. (1995), Fundamentals of artificial neural networks, MIT press.

Hawkins, D. M. (2004), The problem of overfitting, Journal of Chemical Information and Computer Sciences 44, 1, 1–12.

Jones, K. (1977), Some epistemological considerations of paradigm shifts. Basic steps towards a formulated model of alternation, The Sociological Review 25, 2, 253–272.

Kaneyoshi, T. / Jaščur, M. / Fittipaldi, I. (1993), Transverse Ising model with arbitrary spin, Physical Review B 48, 1, 250.

Keil, F. C. (1991), The emergence of theoretical beliefs as constraints on concepts, in: Carey, S. (ed.), The Epigenesis of Mind. Essays on Biology and Cognition, Hillsdale, NJ: Erlbaum, 237–256.

Kurzweil, R. (2012), How to create a mind. The secret of human thought revealed, New York: Penguin.

Mitchell, M. (2009). Complexity. A guided tour, Oxford: Oxford Univ. Press.

Mitchell, M. / Newman, M. (2002), Complex systems theory and evolution, in: Pagel, M. (ed.), Encyclopedia of Evolution, Oxford: Oxford Univ.-Press. https://www.academia.edu/2851794/Complex_systems_theory_and_evolution [03.09.2020].

Pennycook, G. / Cheyne, J. A. / Seli, P. / Koehler, D. J. / Fugelsang, J. A. (2012), Analytic cognitive style predicts religious and paranormal belief, Cognition 123, 3, 335–346.

Rodriguez, N. / Bollen, J. / Ahn, Y. Y. (2016), Collective dynamics of belief evolution under cognitive coherence and social conformity, PloS One 11, 11, e0165910. DOI: 10.1371/journal.pone.0165910.

Ruffman, T. / Perner, J. / Parkin, L. (1999), How parenting style affects false belief understanding, Social Development 8, 3, 395–411.

Sodian, B. / Zaitchik, D. / Carey, S. (1991), Young children's differentiation of hypothetical beliefs from evidence, Child Development 62, 4, 753–766.

Sutton, R. S. / Barto, A. G. (2018), Reinforcement learning. An introduction, Cambridge, Mass.: MIT Press.

Toner, K. / Leary, M. R. / Asher, M. W. / Jongman-Sereno, K. P. (2013), Feeling superior is a bipartisan issue. Extremity (not direction) of political views predicts perceived belief superiority, Psychological Science 24, 12, 2454–2462.