Learning from patterns: information retrieval and visualisation issues between bioimage informatics and digital humanities

Irene Cazzaro[1][0000-0002-9484-1980]

[1]Dipartimento di Architettura, Alma Mater Studiorum Università di Bologna
irene.cazzaro2@unibo.it

Abstract. The large amount of data generated in different fields, among which bioimage informatics and digital humanities, is increasingly requiring appropriate automatic processing techniques, such as computer vision, data mining and particular visualisation tools, to extract information out of complexity and to clearly display it.

This has led, in digital humanities, to the use of pattern recognition techniques similar to those applied in biology, chemistry and medical studies, but where patterns to be analysed and segmented are extracted from texts, images, audio-visual and online media rather than from cells and tissues. Regularities can be recognised through machine learning, based on artificial neural networks that are modelled, to some extent, after the brain's structure, showing a variety of analogies between natural and artificial world.

These processes can also add information to 3D models for cultural heritage: data mining technologies allow information retrieval from archives and repositories, as well as the comparison of data in order to better understand the context of – and relationships between – works of art, thus producing knowledge enhancement.

Various tools to describe complexity are here analysed not only for their educational aim, but also for their heuristic value, allowing new discoveries and connections between different disciplines.

Keywords: Bioimages learning, image-based education, visual simulation and modelling learning, visual-based research methods, visual studies.

1 Introduction

1.1 Artificial neural networks

The extraction of useful information out of complexity, a principle of both bioimage informatics and digital humanities, generates a large amount of data through automatic processing, namely techniques such as computer vision, data mining and particular visualisation tools, able, as Herbert Simon would have said, "to find pattern hidden in apparent chaos" (Simon 1968).

These kinds of operations are increasingly common now, considering that we can easily access big data in many different fields, in both the natural and the artificial domain. One of the most used models to analyse complex data is based on artificial neural networks, which employ machine learning, one of the major branches of artificial intelligence.

When we speak of artificial neural networks (Wang 2003; Bishop 2006), we refer to systems whose aim is to process information simulating, as much as possible, the functioning of biological networks, where interconnected computers play the role of (simplified) neurons acquiring information from the external world, elaborating it and giving an output useful for data analysis and decision-making.

In particular, they are composed of input, hidden and output nodes, each of which is activated if the received signal is higher than an activation threshold and is connected with the others through weighted connections (fig. 1). The number of links (synapses) can also increase or decrease based on the type of stimulus.

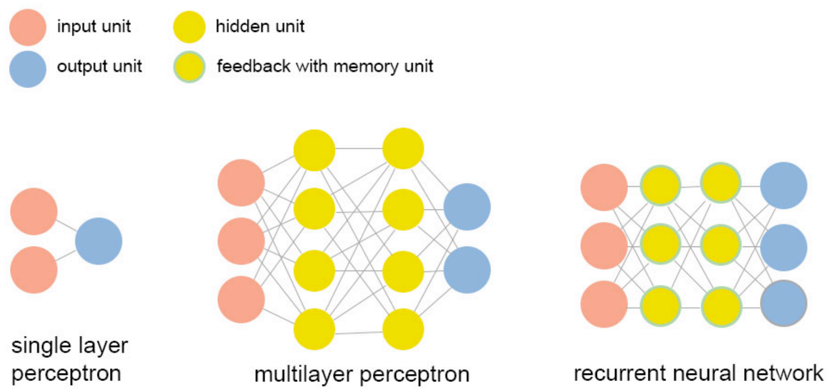

Fig. 1. Three phases of the development of artificial neural networks, from the single layer perceptron (1960s) to the more complex recurrent neural network (1990s), where feedback with a memory unit becomes possible (1990s). Author's editing based on https://www.allerin.com/blog/3-types-of-neural-networks-that-ai-uses (accessed 9 June 2021).

This has become the standard approach to data analysis in science and engineering, but also in many other fields dealing with big data and transforming our way of approaching information.

The development of these technologies has been possible mainly for two reasons: on the one hand, the exponential growth of data collected and accessible online; on the other hand, the development of more powerful computers (especially in relation to GPUs) able to process this large amount of data and to analyse them at a morphological, structural and dynamic level.

1.2 Patterns

The term "pattern" can be used as a synonym of "texture", thus indicating an arrangement of shapes and colours seen physically, for example on the fur of animals or on clothes. More in general, a pattern is a repetition of a sequence in a set of raw data: in the context of this research, we primarily refer to patterns observed mathematically through algorithms and used to identify regularities in language, images or classes of objects (Haken 2004).

Machine learning algorithms use pattern recognition to process data according to statistics and especially using the nearest neighbour search, an automatic process that makes it possible, given a single point in a set, to find its closest or most similar neighbours within a certain distance.

We can see, therefore, that the classification of data can be grounded on already gained knowledge, but also on statistical information automatically extracted from patterns.

The application potential of these processes ranges from speech recognition and speaker identification to multimedia document recognition, automatic biological analysis and medical diagnosis, where raw data are converted in a form that is manageable by a machine, involving classification and clustering (Levenberg, Neilson, and Rheams 2018).

In classification, a class label is assigned to a pattern based on an abstraction generated using a set of training patterns and domain knowledge (supervised learning), whereas clustering generates a partition of data based on the recognition of patterns that helps decision-making (unsupervised learning). Supervised and unsupervised learning are the main machine learning paradigms together with reinforcement learning, which helps a computer to learn appropriate behaviour through repeated "trial-and-error" interactions with a dynamic environment where a decision depends on the current state of a system and determines the following one.

Pattern recognition should recognise familiar patterns quickly and accurately, even if they are seen from different angles or partly hidden. It should also recognise and classify unfamiliar objects.

The process of learning, through which a system is trained and becomes adaptable to accurately give results, depends on which algorithms are used on a dataset that is usually divided in two categories:

a) A training set, such as a series of images used to train the system and build a model;
b) A testing set, to test the system and verify if it is correct.

These are the main features of artificial neural networks and pattern recognition techniques that are used in both natural and artificial domains and that have also shown similarities with other models employed to analyse complexity, as we will see in the next paragraphs.

1.3 Short history

The application of deep learning technologies has been one the most important paradigm shifts in the last years, gaining popularity especially in relation to artificial neural networks.

The first mathematical model of a neuron was proposed in 1943 by Warren McCulloch and Walter Pitts: multiple input data were transformed into a single output according to thresholds (McCulloch and Pitts 1943); the artificial neuron could also be combined with other elements creating a network to solve simple Boolean functions. D. O. Hebb in 1949 proposed the first learning hypotheses based on the brain's complex models (Hebb 1949). In the same years, information theory was proposed (Shannon 1948; Shannon and Weaver 1949), as well as mathematical models to account for complexity (Waddington 1940; Turing 1952).

In 1958 Frank Rosenblatt proposed the "perceptron" (Rosenblatt 1958), the first scheme of neural network that was able to assign weights to properties by learning from examples (fig. 1 left): it is a probabilistic model using pattern recognition. The interest for similar models continued for a decade and had a strong influence in the field of computational geometry (Minsky and Papert 1969), even though many problems could not be solved by the perceptron model.

We have to wait until 1986 for the enhancement of these systems through the error back-propagation algorithm (Rumelhart, Hinton, and Williams 1986), that generalises the learning algorithm of the perceptron, and until the 1990s (Schmidhuber 1992) for the recurrent and multi-layer neural networks (fig. 1 right), which have been applied to image and 3D object recognition.

Artificial neural networks, however, were hard to train and they underwent a period of relative crisis until, between 2009 and 2012, computational neural networks began winning prizes in competitions, approaching human performance in many tasks (Graves 2012; Krizhevsky, Sutskever, and Hinton 2017): from that moment on, machine learning has gained ground in many different fields.

2 Some applications in different domains

2.1 Biology, chemistry, medical studies

Techniques like nearest neighbour search and segmentation are widely used for pattern recognition in biology, chemistry and medical studies (Coelho et al. 2010; Meijering 2020), where they are often combined with visualisation in false colours to better identify the segmented areas (tissues, cells, proteins…).

Microscopy and image analysis are used for both quality and quantification purposes in complex model systems, involving processes like the isolation of cells and tissues at a single-cell resolution and the labelling of different phenomena.

Biological algorithms are designed in order to allow the following operations:

a) Image restoration and pre-processing, with the aim of starting with good data. The most frequent operations in this phase are cropping (to reduce the area of interest), inversion (the creation of negative images), filtering (to extract significant data), colour extraction, illumination correction, elimination of blurring;

b) Segmentation, that allows the identification of objects: single entities are measured, individual cells and structures are identified (object detection), thresholds are fixed based on the frequency of pixel values, shapes that overlap each other are divided (untangling);
c) Tracking, i.e. observing the movement of objects over time;
d) Object classification, where relevant properties are identified;
e) Quantification, where relevant properties are quantified;
f) Visualisation, often using false colours, to better identify the segmented areas.

The development of these processes has surely been influenced by research on the brain, but even the opposite has happened: models such as "connectionism" (Rumelhart et al. 1968), applied in cognitive science, use artificial neural networks to explain the functioning of the brain: the analogy, rather than between mind and computer, is between natural and artificial neural networks distributing the various activities through connections and computational units, in a behavioural framework based on the connection between stimulus and response.

2.2 Glass: from liquid to solid "amorphous" state

Artificial neural networks are also giving an impulse on studies trying to identify the patterns of transformation of glass from liquid to solid "amorphous" state (Bapst et al. 2020), where the molecules remain in a seemingly disordered state, much like a liquid. It is therefore fundamental to understand how atomic-scale properties define the visible features of many solid materials.

The physical aspect of glass is predicted with the aid of a structure with nodes representing particles and edges representing interactions. Each particle in the simulation travels a distance and, after several iterations, a graph network takes shape.

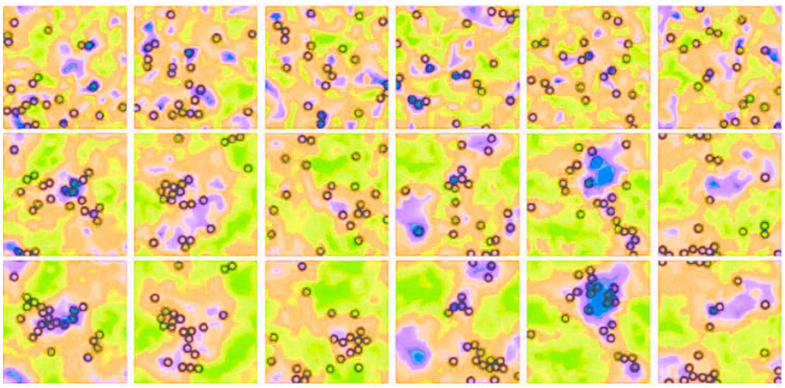

Fig. 2. Artificial neural networks used to study the passage of glass from liquid to solid "amorphous" state. The coloured spots represent the results obtained with artificial neural networks, whereas the circles represent the actual behaviour of glass molecules. The model is accurate when circles are close to the purple or blue spots. Author's editing based on (Bapst et al. 2020).

Several datasets are constructed corresponding to predictions of mobility for different time spans and temperatures. These have proved to be quite reliable compared to reality, even though the accuracy decreases when time spans increase (fig. 2).

2.3 Digital Humanities

Similar methods are used in digital humanities (Warwick, Terras, and Nyhan 2012), where patterns to be analysed and segmented are extracted from texts or, sometimes, from images, audio-visual and online media. Regularities can be recognised through machine learning, based on artificial neural networks, showing a variety of analogies between natural and artificial world, as happens in many studies focusing on networks (Buchanan 2004).

Data and text mining allows the discovery of patterns in large datasets from where information is extracted with methods bound to machine learning, statistics and database systems. This constitutes the analysis step of the “knowledge discovery in databases” process (Frawley, Piatetsky-Shapiro, and Matheurs 1992). It doesn’t only involve the act of extracting knowledge, but also – similarly to bioimage analysis – database and data management processing, model and inference considerations, segmentation, classification, post-processing, visualisation and online updating.

The scope of this discipline is really wide if we think of information that can be extracted from digital archives and that can be connected through standardised technologies involving linked open data in the “semantic web” context (Berners-Lee, Hendler, and Lassila 2001).

2.4 Digital Heritage Studies

These processes can also add information to 3D models for cultural heritage: data mining technologies allow information retrieval from archives and repositories, as well as the comparison of data in order to better understand the context of – and relationships between – works of art, thus producing knowledge enhancement: the challenging opportunity represented by automatic visual retrieval makes it possible to match images based not on standardised keywords as happens with linked open data, but on visual information, a task whose “complexity calls for a truly interdisciplinary endeavor” (Bell and Ommer 2016).

Segmentation can be done by identifying patterns in 3D models, both automatically and with the help of the human eye, thus classifying the different components hierarchically according to “ontologies”, which can be used to enhance interoperability in virtual research environments (Statham 2019; Champion and Rahaman 2020). Visualisation devices become then relevant, as well as in biology, in dealing with online platforms where the shared models have to comply with principles mainly deriving from the Gestalt theory and the “semiology of graphics” (Bertin 1967; Tufte 1990). As an example, when we refer to hypothetical reconstructions of unbuilt or destroyed artefacts, visualisation issues arise right in the use of false colours to illustrate the segmentation in different temporal phases or levels of uncertainty.

The aim of these tools is surely educational, but also heuristic, allowing new discoveries and connections between data coming from different archives or virtual research environment that can be easily shared online. As an example, machine learning is already used by *Google Arts and Culture*, with the aim of protecting our past and cultural heritage[1].

3 Modelling complex systems

Machine learning is a field in constant expansion and we have seen that many different phenomena can be analysed through neural networks: the behaviour of granular materials, biological systems, social phenomena, texts, images and cultural heritage. However, it is not the only model capable of that, even though now it is the most successful one.

The examples that we have seen are based on local constraints where the position of some elements inhibits the motion of others, in a complex and cooperative dynamics: this is similar to the mechanism of reaction-diffusion models, basically composed of two substances, an activator and an inhibitor, that react between them and diffuse according to the concentration of the same substances in the neighbourhoods. An example of this is Alan Turing's morphogenetic model (Turing 1952), where the substances are called "morphogens" and are considered the responsible not only for the differentiation of an embryo into a living being, but also for the appearance of an actual pigmentation pattern on animals. The model has also been used to account for the development and interaction between cities (Pumain 1998; Allen and Sanglier 2010), as well as the behaviour of ants or the formation of sand dunes (Ball 2015).

This ability of self-organisation is also shared with multi-agent (or self-organising) systems (Wooldridge 2009), which are computerised systems composed of many interacting software agents able to perform particular tasks. Their application to machine learning has also given rise to "agent mining" (Cao, Gorodetsky, and Mitkas 2009). The goal is to search for the explanation of a complex phenomenon (such as online trading or social structure modelling) into the collective behaviour of these agents that act according to simple rules.

These models are also close to cellular automata (Wolfram 2002) and L-systems (Prusinkiewicz and Lindenmayer 1990), where the iteration of a simple rule through time is studied. The rule only depends on the value of a parameter assigned to an element compared to the values of its neighbours, something that reminds us the nearest neighbour search, but also the "diffusion" phase in reaction-diffusion models.

The shared framework of these models also involves information theory (Shannon 1948; Shannon and Weaver 1949), network theory (Buchanan 2004) and chaos theory (Lorenz 1963), helping us identify the similarities (fig. 3) between these attempts to describe complex systems with models that remain simplifications of reality, but still they can be useful (Box 1976).

---

1 Some of the experiments performed by *Google Arts and Culture* in relation to the use of artificial intelligence for cultural purposes can be found on: https://artsandculture.google.com/story/unlock-culture-at-home-with-machine-learning/kwKSLHCd3edAIg (accessed 9 June 2021).

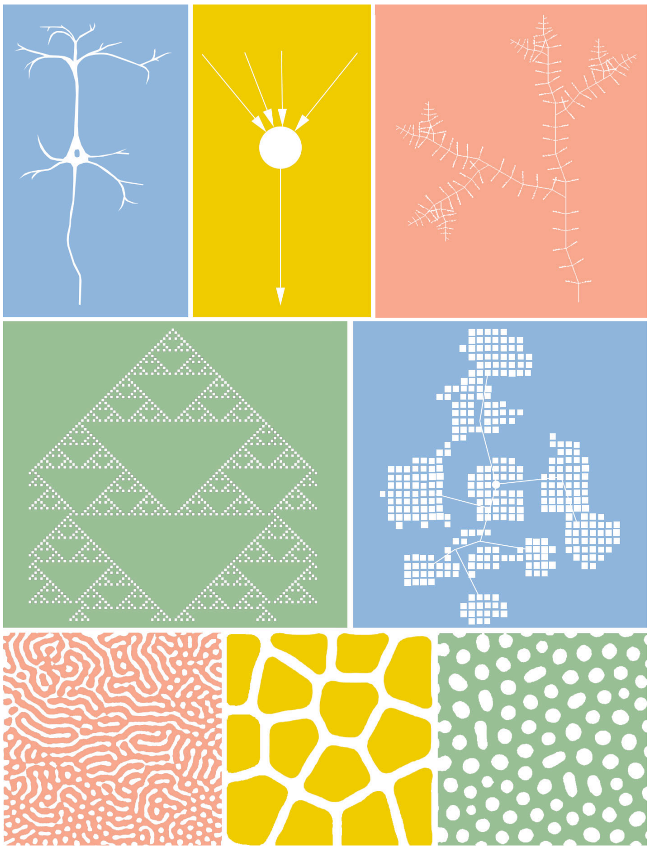

Fig. 3. Visual and conceptual analogies between different phenomena that have been described making use of more or less complex algorithms. From top left to bottom right: the representation of a neuron, an artificial neural network scheme, an L-system simulating the growth of a plant, a Wolfram elementary cellular automaton, an urban growth simulation, three different Turing patterns. Author's editing based on (Turing 1952; Wolfram 2002; Raimbault, Banos, and Doursat 2016), http://paulbourke.net/fractals/lsys/ (accessed 9 June 2021).

4 Conclusions

In this short overview, we have seen a wide range of models and applications that have not only a primary role in the educational field, but also an incredible heuristic power in extracting data out of the complexity of the real world and classifying them, thus allowing new discoveries and helping manage large amounts of data. This is the reason why these tools are widely employed by giants of the web such as Google, Facebook, Apple, etc. who have heavily invested in machine learning to label and extract information from images, texts and multimedia.

Machine learning has shown its advantages not only in quantitative predictions, but also in qualitative understanding. Reality is undoubtedly far more complex than the most up-to-date artificial neural network, which still requires – and will probably continue to require – the presence of man. We might therefore conclude that these models remain a way to augment, rather than replace, human understanding.

References

Allen, Peter M., and Michèle Sanglier. 2010. "A Dynamic Model of Growth in a Central Place System." *Geographical Analysis* 11 (3): 256–72.

Ball, Philip. 2015. "Forging Patterns and Making Waves from Biology to Geology: A Commentary on Turing (1952) 'The Chemical Basis of Morphogenesis.'" *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 370 (1666).

Bapst, V., T. Keck, A. Grabska-Barwińska, C. Donner, E. D. Cubuk, S. S. Schoenholz, A. Obika, et al. 2020. "Unveiling the Predictive Power of Static Structure in Glassy Systems." *Nature Physics* 16 (4): 448–54.

Bell, Peter, and Björn Ommer. 2016. "Digital Connoisseur? How Computer Vision Supports Art History." In *Il Metodo Del Conoscitore: Approcci, Limiti, Prospettive*, 1–15. Roma: Editoriale Artemide.

Berners-Lee, Tim, James Hendler, and Ora Lassila. 2001. "The Semantic Web: A New Form of Web Content That Is Meaningful to Computers Will Unleash a Revolution of New Possibilities." *Scientific American* 284 (5): 1–5.

Bertin, Jacques. 1967. *Semiology of Graphics: Diagrams, Networks, Maps*. Redlands, California: Esri Press.

Bishop, Christopher M. 2006. *Pattern Recognition and Machine Learning*. New York: Springer-Verlag.

Box, George E. P. 1976. "Science and Statistics." *Journal of the American Statistical Association* 71 (356): 791.

Buchanan, Mark. 2004. *Nexus. Perché la natura, la società, l'economia, la comunicazione funzionano allo stesso modo*. Translated by L. Serra. Milano: Mondadori.

Cao, Longbing, Vladimir Gorodetsky, and Pericles A. Mitkas. 2009. "Agent Mining: The Synergy of Agents and Data Mining." *IEEE Intelligent Systems* 24 (3): 64–72.

Champion, Erik, and Hafizur Rahaman. 2020. "Survey of 3D Digital Heritage Repositories and Platforms." *Virtual Archaeology Review* 11 (23): 1–15.

Coelho, Luis Pedro, Estelle Glory-Afshar, Joshua Kangas, Shannon Quinn, Aabid Shariff, and Robert F. Murphy. 2010. "Principles of Bioimage Informatics: Focus on Machine Learning of Cell Patterns." *Lecture Notes in Computer Science,* 6004 LNBI: 8–18.

Frawley, William J., Gregory Piatetsky-Shapiro, and Christopher J. Matheurs. 1992. "Knowledge Discovery in Databases: An Overview." *AI Magazine* 13 (3): 57–70.

Graves, Alex. 2012. *Supervised Sequence Labelling with Recurrent Neural Networks*. Berlin, Heidelberg: Springer-Verlag.

Haken, Hermann. 2004. "What Are Patterns?" *Synergetic Computers and Cognition. Springer Series in Synergetics* 50: 9–17.

Hebb, Donald O. 1949. *The Organization of Behavior: A Neuropsychological Theory*. New York: John Wiley & Son Ltd.

Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton. 2017. "ImageNet Classification with

Deep Convolutional Neural Networks." *Communications of the ACM* 60 (6): 84–90.

Levenberg, Lewis, Tai Neilson, and David Rheams. 2018. *Research Methods for the Digital Humanities*. London: Palgrave Macmillan.

Lorenz, Edward N. 1963. "Deterministic Nonperiodic Flow." *Journal of the Atmospheric Sciences* 20: 130–41.

McCulloch, Warren S., and Walter H. Pitts. 1943. "A Logical Calculus of the Ideas Immanent in Nervous Activity." *Bulletin of Mathematical Biophysics* 5: 115–33.

Meijering, Erik. 2020. "A Bird's-Eye View of Deep Learning in Bioimage Analysis." *Computational and Structural Biotechnology Journal* 18: 2312–25.

Minsky, Marvin L., and Seymour A. Papert. 1969. *Perceptrons: An Introduction to Computational Geometry*. Cambridge: MIT Press.

Prusinkiewicz, Przemyslaw, and Aristid Lindenmayer. 1990. *The Algorithmic Beauty of Plants*. New York: Springer-Verlag.

Pumain, Denise. 1998. "Les modèles d'auto-organisation et le changement urbain". *Cahiers de géographie du Québec*, 42(117), 349–366.

Raimbault, Juste, Arnaud Banos, and René Doursat. 2016. "A Hybrid Network/Grid Model of Urban Morphogenesis and Optimization," no. February.

Rosenblatt, Frank. 1958. "The Perceptron: A Probabilistic Model for Information Storage and Organization in the Brain." *Psychological Review* 65 (6): 386–408.

Rumelhart, David E., Geoffrey E. Hinton, and Ronald J. Williams. 1986. "Learning Representations by Back-Propagating Errors." *Nature* 323 (6088): 533–36.

Rumelhart, David E., James L. Mc Clelland, and CORPORATE PDP Research Group. 1968. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Vol. 1: Foundations*. Vol. 1. Cambridge: MIT Press.

Schmidhuber, Jürgen. 1992. "Learning Complex, Extended Sequences Using the Principle of History Compression." *Neural Computation* 4 (2): 234–42.

Shannon, Claude E. 1948. "A Mathematical Theory of Communication." *The Bell System Technical Journal* 27 (3): 379–423.

Shannon, Claude E., and Warren Weaver. 1949. *The Mathematical Theory of Communication*. Urbana and Chicago: University of Illinois Press.

Simon, Herbert A. 1968. *The Sciences of the Artificial*. Cambridge, M.I.T. Press, C.

Statham, Nataska. 2019. "Scientific Rigour of Online Platforms for 3D Visualisation of Heritage." *Virtual Archaeology Review* 10 (20): 1–16.

Tufte, Edward R. 1990. *Envisioning Information*. Cheshire, Connecticut: Graphic Press.

Turing, Alan Mathison. 1952. "The Chemical Basis of Morphogenesis." *Philosophical Transactions of the Royal Society of London. Series B, Biological Sc.* 237 (641): 37–72.

Waddington, Conrad Hal. 1940. *Organisers & Genes*. Cambridge [England: University Press.

Wang, Sun-Chong. 2003. "Artificial Neural Network." In *Interdisciplinary Computing in Java Programming*, 743:81–100. Boston, MA: Springer.

Warwick, Claire, Melissa M Terras, and Julianne Nyhan. 2012. *Digital Humanities in Practice*. London: Facet.

Wolfram, Stephen. 2002. *A new kind of science*. Champaign, Ill.: Wolfram Media.

Wooldridge, Michael. 2009. *An Introduction to MultiAgent Systems*. Chichester, West Sussex: John Wiley & Sons Ltd.