A Cybernetics Perspective on Data Science: Macro and Micro Views

Cyril S. KU
Department of Computer Science
William Paterson University
Wayne, NJ 07470, USA

Joseph R. LARACY
Department of Systematic Theology
Department of Mathematics and Computer Science
Seton Hall University
South Orange, NJ 07079, USA

Thomas J. MARLOWE
Department of Mathematics and Computer Science
Seton Hall University
South Orange, NJ 07079, USA

Jin-A CHOI
Department of Communication
William Paterson University
Wayne, NJ 07470, USA

ABSTRACT

In this paper, data science is considered from a cybernetic perspective in two viewpoints. After a brief review of cybernetics, a partial conceptual view of a data science framework is provided. Several layers are identified, working from the software engineering life cycle macro perspective of a data analysis system, through production of a machine learning model to mine knowledge and predict business and product trends, to the micro perspective of a specific analysis, in this case using an artificial neural network. How the layers fit, individually and collectively, into a cybernetic system, identifying feedback loops and their interactions are described. Finally, the advantages and disadvantages of understanding the modern data science life cycle from the cybernetics perspective, and insights to be gained from this perspective are discussed.

Keywords: Artificial Neural Networks, CRISP-DM Life Cycle, Cybernetics, Data Science, Machine Learning.

1. INTRODUCTION

Models that unify distinct phenomena can provide insights, and may lead to better understanding, ease integration or management of tools and processes, or allow improvements and optimizations. In this paper, data science is reconsidered in the light of a cybernetic model, in order to improve understanding, and with a goal of leading to enhancements of the data science process.

"Cybernetics" comes from the Greek term, κυβερνήτης (kybernētēs), which denotes a pilot, governor, steersman, or rudder. It is the study of communication and control in systems (biological, engineered, and socio-technical), often characterized by feedback, time delays, and nonlinearities. The simplest such systems consist of a controller and a subsystem carrying out some function, whose state and outputs are observed by the controller, which uses that information in determining which control actions to take (or control signals to send).

A thermostat is a classic example of a simple cybernetic system. In more complex systems, inputs to the controller may include environmental factors, and may have configurable parameters—in the thermostat example, humidity could be an environmental factor, and a human-adjustable target temperature a configurable goal parameter. Second-order cybernetics also considers human controllers, who may be influenced by the controlled system, and are open to reflection (on the state of the system and environment) and reflexion (on internal state), so that the controller/observer becomes part of the process.

Data science as a discipline uses mathematical, statistical, and computing techniques, many derived from artificial intelligence (AI), to analyze often massive sets of environmental, user-generated, and/or experimental data. Individual data analytic models and techniques are parameterized, both by user/developer-adjustable hyperparameters (parameters at the higher/macro level), and by internal parameters (parameters at the lower/micro level), which in many AI-driven applications the tool modifies for improved results. AI-driven applications include neural networks, genetic algorithms, and simulated annealing with heuristics.

Another dimension lies in textual and graphical visualizations of results, where displays and overall layout should themselves be tuned to optimize utility, comfort, and aesthetics. This interacts with user experience (UX) since what is important is to convey information. Visualizations and displays will be modified both for technical considerations and by user feedback, or possibly by experimentation with alternatives, using hypothesis-driven design (HDD) [1].

There are therefore several processes involved in a data science-based application: choosing analyses, integrating results, tuning hyperparameters, carrying out analyses, and selecting display and layout of the results. Each of these can be viewed as a cybernetic system, and in addition, development of the system can itself be viewed as a cybernetic system, especially if an incremental iterative process such as agile development [2] is used.

The rest of this paper looks further, first at cybernetic systems, and then at data science. It then assesses facets of data science from a cybernetic perspective, identifying feedback loops, and considering interaction of those loops. Finally, in the conclusions, it considers implications of these findings, and looks at possible future directions.

2. CYBERNETIC SYSTEMS

The birth of first-order cybernetics occurred in the United States and the United Kingdom in the aftermath of the Second World War. It is perhaps best understood as a transdisciplinary approach with foundations in dynamical systems theory, logic and discrete mathematical modeling, electrical and computer engineering, mechanical engineering, and neuroscience. The MIT professor of

ISBN: 978-1-950492-64-0
ISSN: 2771-0947

https://doi.org/10.54808/WMSCI2022.01.47

mathematics, Norbert Wiener, introduced the term "cybernetics" in his 1948 book on the study of control and communication in the animal and the machine [3].

Wiener's ideas and insights contributed substantially to the emerging fields of neuroscience, analog computing, artificial intelligence, control theory, and communication theory. In addition, his mathematical work in probability theory provided the foundations for Claude Shannon to develop modern information theory [4].

The 1956 book, *An Introduction to Cybernetics*, by W. Ross Ashby [5], a psychiatrist, further advanced the emerging cybernetics paradigm. Ashby is well known for formulating the crucial cybernetic notions of

- the law of requisite variety (i.e., the greater the variety of actions available to a control system, the greater the variety of perturbations it is able to manage),
- the principle of self-organization (i.e., a deterministic dynamic system evolves towards a state of equilibrium that can be described in terms of a basin of attraction of surrounding states), and
- the principle of regulatory models (i.e., every good regulator of a system must be, or contain a model of that system).

Ashby is also the inventor of the homeostat, an early cybernetic electro-mechanical system capable of adapting itself to its environment.

During his 1949 lectures at the University of Illinois, John von Neumann shared his quest to design a system whose complexity could develop in a way analogous to living organisms subject to Charles Darwin's notion of natural selection. Drawing on his expertise in physics, mathematics, and computer science, von Neumann explicated the logical requirements for machine self-replication in a cellular automaton environment, thus articulating the notion of the universal constructor. The 1966 book, *Theory of Self-Reproducing Automata*, completed by Arthur W. Burks after von Neumann's death, not only contributed to automata theory, but also to the study of "artificial life" and complex systems theory [6].

FOCS (First-Order Cybernetic System)

In a first-order system, a subsystem receives inputs from its environment as well as from the controller. Control actions are based on the estimated state of the subsystem. The state estimation process is guided by system outputs (which in turn are inputs to the controller) and pre-programmed algorithms in the controller. This allows the subsystem to maintain an equilibrium or evolve toward another goal state. In either case, the subsystem should produce the intended external effects envisioned by the system designer. See Figure 1 below.

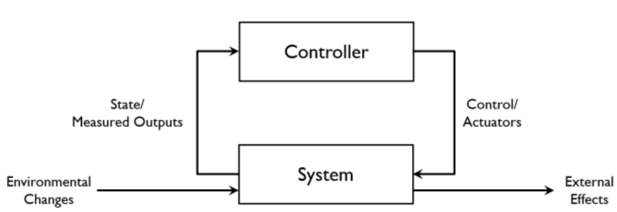

Figure 1: FOCS

The "bridge" figure between first-order and second-order cybernetics was Heinz von Foerster, a University of Illinois professor of biophysics and electrical engineering. Von Foerster founded the Illinois Biological Computer Laboratory (BCL) in 1958. Over the next sixteen years, von Foerster and his colleagues carried out research in the areas of:

- bionics (i.e., the application of biological methods and systems found in nature to the study and design of engineering systems),
- bio-inspired computing (i.e., analyzing, formalizing, and implementing biological processes using computers), and
- self-organizing systems. See [7].

According to Von Foerster, second-order cybernetics emerged as a consequence of the BCL research to develop a model of the human mind [8]:

> [A] brain is required to write a theory of a brain. From this follows that a theory of the brain, that has any aspirations for completeness, has to account for the writing of this theory. And even more fascinating, the writer of this theory has to account for her or himself. Translated into the domain of cybernetics; the cybernetician, by entering his own domain, has to account for his or her own activity. Cybernetics then becomes cybernetics of cybernetics, or second-order cybernetics.

The new "cybernetics of observing systems" shifted the disciplinary focus to the social sciences, communication, and the humanities [9, 10, 11]. In particular, cyberneticists began to consider issues related to cognitive science, epistemology, and the philosophy of science, with clear mutual influence.

SOCS (Second-Order Cybernetic System)

Second-order cybernetics is often described as the reflexive practice of cybernetics. The cyberneticist and other humans in the loop are part of the system under study. They cannot be "abstracted" from the system, e.g., a command-and-control system for ballistic missile defense which includes human, information, and physical resources.

Key features of second-order systems are reflection (on the state of the system and environment) and reflexion (on internal state) by the observer/controller. The observer/controller must be capable of

- judgment in application of rules (similar to a first-order system),
- reflection leading to dynamic control or rule modification, and
- reflexion leading to changes in observer internal state.

Generally speaking, the observed system is capable of broader interaction with the observer.

As Figure 2 below indicates, the observer/controller receives inputs not only from the subsystem but also directly from the external environment and the internal state of the subsystem. In addition, the observer/controller is self-reflective. The observer/controller may also have direct impact on the external effects of the system. Finally, the subsystem is capable of self-reaction, independent of the observer/controller.

In 2021, Thomas J. Marlowe et al. proposed a transitional-order cybernetic paradigm [12]. A transitional-order system has many characteristics of a second-order system. However, the observer/controller has only limited capacity for reflection and none for reflexion. Examples include a system with an artificially intelligent, self-modifying, and learning-capable observer/ controller.

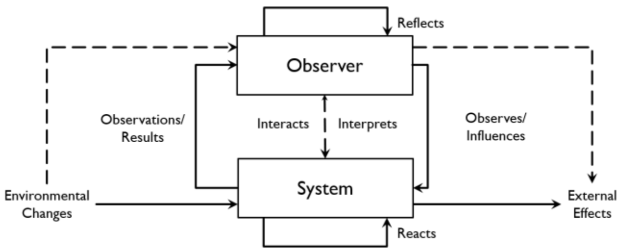

Figure 2: SOCS

3. DATA SCIENCE

Data science is an interdisciplinary field to extract knowledge from very large data set (often called big data). Its techniques, approaches, and algorithms come primarily from statistics, discrete mathematics and operations research, and computer science, particularly artificial intelligence, while its applications span the gamut, from, for example, business analytics and economics, through social science and politics, to genetics and medicine. The goal of data science, particularly in business applications, is to improve decision making by basing decisions on insights from the big data, or, as in genetics, to enhance or extend knowledge and understanding of the discipline [13, 14, 15].

One important facet of data analysis is data mining. Data mining looks for patterns—typically unexpected—in various form of data (in a structured form such as a record in a relational database table, or an unstructured form such as text, graphics, or molecular structures). The patterns detected by data mining algorithms represent the knowledge, trends, predictions, anomaly detection of the big data.

Machine learning is one of the techniques used for data mining, deriving mostly from artificial intelligence. Its techniques include association, classification, clustering, outlier/anomaly detection, neural networks, pattern recognition, genetic algorithms, and multivariable statistical analyses such as principal components algorithms [16, 17, 18].

4. MACRO VIEW: CRISP-DM LIFE CYCLE

To produce a software product, an organization follows a specific process or software engineering model to develop the software. These processes, of which agile development processes [19] are examples, follow the SDLC (Software Development Life Cycle) [20]. Analogously, the CRISP-DM (Cross Industry Standard Process for Data Mining) is widely used in data science projects and specifically for producing and tuning data mining and machine learning models [13]. CRISP-DM is also known as the data science life cycle.

Figure 3 shows the CRISP-DM life cycle for a data science project. The first two stages involve business understanding and data understanding where the data scientist is trying to define the goals of the project by understanding the business needs and the data that the business needs for the project. Most of data comes from the organization internally and some data may need to be obtained externally. This leads to the data preparation stage. Data preparation consumes as much as 70% to 80% of the effort for the whole life cycle [13]. Data preparation includes data extraction, data cleaning, data transformation, data migration and integration, data normalization, data aggregation, and data loading. One of the challenges is the get the relevant data to build a data mining or a machine learning model. After the modeling stage, the data scientist as well as other users who will use the model need to evaluate the effectiveness or correctness of the model. The evaluation results may lead the data scientist to restart the life cycle if the model is not effective. If the results turned out to be feasible or useful, then it will be deployed to be used by the organization. With the passage of time, even well formulated models eventually need to be updated. The block arrows surrounding the diagram in Figure 3 indicate the iterative nature of the CRISP-DM life cycle.

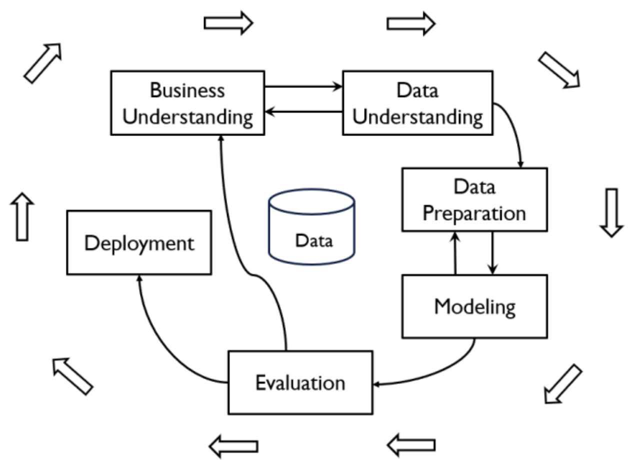

Figure 3: CRISP-DM Life Cycle

Figure 3 describes the overall process of creating a data mining or machine learning model. This will be referred to the macro view of this model building process.

In the following subsection, this macro view of the CRISP-DM process, a second order cybernetics system, is discussed.

The Data Science Life Cycle as A Cybernetic System

As mentioned in Section 2, a second order cybernetics system features feedback loops by the controller and the observer as well as the environment have direct impact on the behavior of the system. The observers are the data scientists and the users who evaluate the modeling process and effectiveness and usefulness of the resulting model. There are many feedback loops in the macro view. The double arrows between business understanding and data understanding indicate the interaction, back and forth between these two stages, and likewise the double arrows between data preparation and modeling. In each case, there may be many iterations between these two stages. As mentioned before, this is the most time-consuming part of the life cycle, as many experiments may need to be performed to select relevant data and come up with a useful model. Third, the users of the resulting model act as the observer of this model building process and will evaluate its effectiveness. A big feedback loop occurs when the observer chooses to restart the whole system. Finally, even though a model is deemed to be effective and useful, the internal evaluation of the model may not reflect reality when it is deployed in the real world. Deployment acts as the environment that influences the effectiveness of the model. The influence includes (1) UX and feedback and (2) performance evaluation.

Figure 4 shows the macro view of the data science life cycle by adding these two environmental factors (see the dashed arrows in the diagram)—the reflection of the cybernetic system. The whole life cycle is a second order cybernetic system. Each feedback loop can be viewed as a mini cybernetic system within the macro view. As described in Section 2, this is cybernetics of cybernetics. That is, it forms a second order cybernetic system. Within each stage, there may be process(es) within itself that can be viewed as a cybernetics system (reflexion of internal state). These cybernetic systems inside the macro view are the micro views of the data science life cycle from the cybernetics perspective.

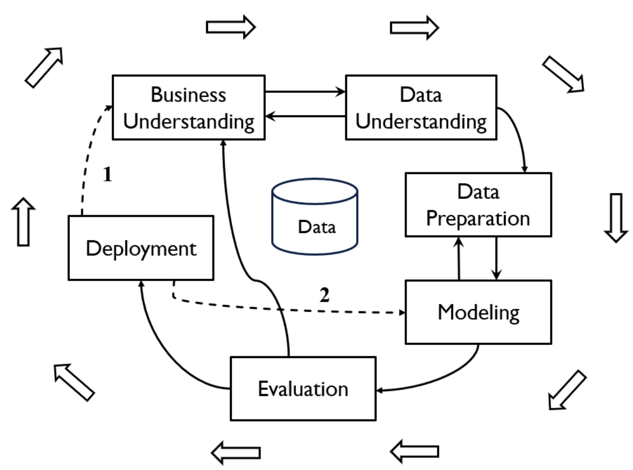

Figure 4: CRISP-DM as a SOCS

5. MICRO VIEW: ARTIFICIAL NEURAL NETWORKS

There are many cybernetic systems (micro views) within a cybernetic system (macro view). The goal of the CRISP-DM process is to produce a data mining or machine learning model. In looking at the resulting model itself, an artificial neural network was chosen to illustrate the micro view of cybernetics. H. Cruse's book [21] describes extensively the modeling of neural networks as cybernetic systems. The artificial neural network is presented below to illustrate the micro view of cybernetics within the macro view.

An artificial neural network (ANN) or simply "neural network" is a machine learning technique. Although it is still unclear as to how the human brain works, it is well known that the brain consists of billions of brain cells called neurons. Scientists have a pretty good idea about the functions of a single neuron. The mystery is how these interconnected neurons form all the functions of the brain. Artificial intelligence uses the mathematical model of a neuron (artificial neuron) and its interconnections to form an ANN. Through the cybernetic mechanism of ANN, an ANN becomes a machine learning algorithm. Just like a biological neuron, the power of the ANN comes from the interconnections of many artificial neurons. Figure 5 shows the three major components of an ANN.

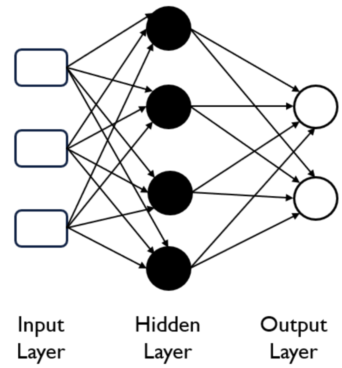

Figure 5: An Artificial Neural Network

A simple ANN consists of three major layers: an input layer, a hidden layer, and an output layer. The diagram shows the network consists of four neurons. Each neuron has its input and output. All the inputs and outputs of all the neurons form the input and output layers respectively. Depending on the application and the design of the algorithm, a user specifies the input and the output if it is a supervised learning model (e.g., classification application). A user specifies the input but lets the ANN generate the output if it is an unsupervised learning algorithm (e.g., clustering application). Consider a simple classification application: the user feeds many pictures of cats and dogs as inputs. The user trains the algorithm by defining the characteristic attributes for cats and for dogs. The output will classify a picture as either a cat or a dog as class labels. After the algorithm learns from the training data, a new picture of a cat fed into the model will have an output saying that the picture is a cat. Using the same example of cats and dogs for an unsupervised application—the user will input many cat and dog pictures into the network. The algorithm will not be able to say a picture is a cat or a dog since unsupervised learning does not use class labels. But a clustering algorithm will be able to put all the cat pictures in a cluster and all the dog pictures in another cluster (if the algorithms has been well trained, of course) based on the input characteristics. In other words, it distinguishes clusters with common characteristics but without names (class labels).

An Artificial Neuron as a Cybernetic System

Let us look at the cybernetic mechanism of a neuron. The mathematical model of a neuron is shown in Figure 6 below [22, 23].

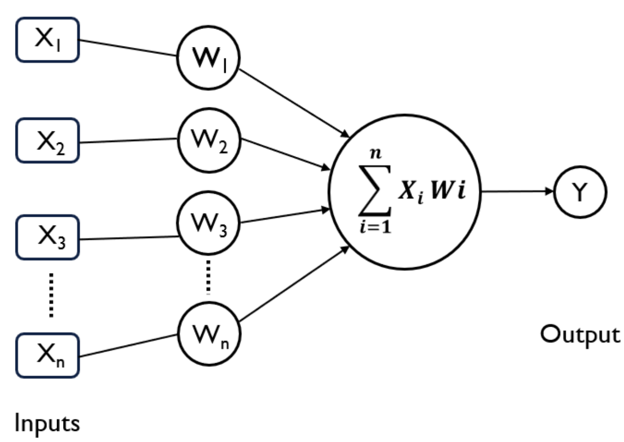

Figure 6: A Mathematical Model of a Neuron

A neuron can have many inputs (X_1 *to* X_n). Each input has a weight (W_i, $1 \leq i \leq n$) associates with it. The processor of the neuron is specified by:

$$X = X_1W_1 + X_2W_2 + X_3W_3 + \ldots + X_nW_n$$

The output of the neuron (Y) is specified by the following activation function:

$Y = +1$ if $X \geq \theta$ or $Y = -1$ if $X < \theta$, where θ is a threshold defined by the user.

If $Y = +1$, the neuron fires and the output will become one of the inputs for another neuron. If $Y = -1$, the neuron will not fire, and it will not influence the neuron that connects to it. Whether a neuron fires or not depends on the inputs and their weight. As the neural network learns from the training data, the weight of each input is adjusted (reinforcement learning). This is the cybernetic mechanism where a feedback loop (back propagation) changes the internal parameters W_i. All these feedback loops were influenced by the hyperparameters of the whole network where the model is being trained. The hyperparameters changed because the whole network is being improved or modified due to evaluation of the results.

Each of the feedback loop strengthens reinforcement learning for each neuron. A layer of neurons strengthens the network. It is not hard to imagine that the more layers of neurons, the more accurate or more patterns can be identified. In other words, more complex problems can be solved. An ANN with many hidden layers, as shown in Figure 7, is called a deep neural network (DNN) or deep learning neural network [24]. Although DNNs are typically feedforward networks (except for recurrent neural networks) in which data flows from input layer to different level of hidden layers and eventually to the output layer, the weights of the neurons are adjusted in response to the evaluation of the results. As demonstrated in the discussion on the workings of a single neuron, DNN is still a model with hyperparameters that will have influences on the internal parameters of different layers of neurons. Again, cybernetic mechanism is exhibited in DNN.

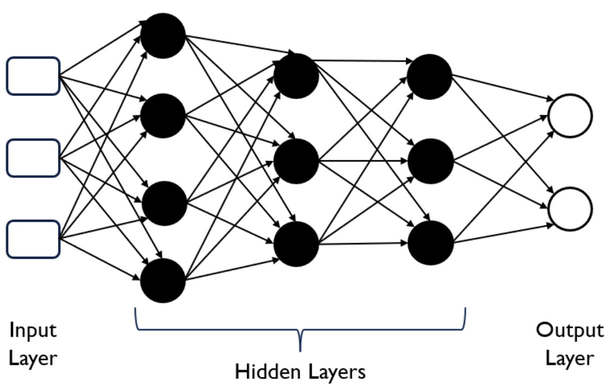

Figure 7: A Deep Neural Network

6. INTERACTIONS BETWEEN MACRO AND MICRO VIEWS

In this paper, the macro view of the data science life cycle from the cybernetics perspective is discussed. Within this macro view, there are many cybernetic micro views. In addition to tools and analyses such as neural nets, genetic algorithms, and multivariable statistical analyses (each of which may have multiple instances in a given application), there are also pre-processing micro-views associated with data acquisition, cleaning, and evaluation, and "front-end" micro views related to visualizations, configurability, and UX. Additional views may be related to extra-functional concerns, including but not limited to quality control, security and privacy, and social and legal issues.

Depending on the scope of the system, the scope and complexity of the goals of the application, and data issues including volume, velocity, variety, veracity, and value (the "5 V's") [25], it may be useful to identify collections of micro views as themselves constituting macro or intermediate level views—either horizontally (data-analysis-display-UX), or vertically, if the system requires addressing multiple challenges or more-or-less independently dealing with multiple data sets.

Each of these views (i.e., micro, intermediate, and macro) will have its own feedback loop(s), and these loops will interact with one another. UX will feed back into display, quality control will feed back not only into the life cycle but also into selection and tuning of individual tools and analyses, and analyses will interact—a statistical principal components analysis will be affected by changes in data evaluation and filtering, and can affect neural network clustering algorithms, which in turn may affect display and visualization—leading to different feedback on the UX and displays.

7. CONCLUSIONS AND FUTURE WORK

Data science encompasses data collection, transformation, modeling, data mining, and data visualization. It is an interdisciplinary field driven by various mathematical models and statistics, databases and data warehouses, data (structured and unstructured) mining, machine learning and deep learning, and artificial intelligence, as well as high performance computing. It is a rapidly growing field in this information age, where decisions are often driven by data. This paper considers the modern discipline of data science from the traditional system science perspective of cybernetics, both at the macro level in application development, and at a micro level, in the use of artificial neural networks. Past efforts that relate data science with systems theory include Berry et al.'s work synthesizing ideas from the mathematical theory of dynamical systems with learning theory to formulate data-driven models of complex systems [26].

The cybernetic perspective allows a unified approach for identifying feedback loops internal to a micro view, and the second-level, interaction feedback loops in intermediate and macro views. Understanding these dependencies allows for better planning and may help to prevent costly errors, delays, and rework in implementing the application. Section 4, for example, shows how applying this perspective allows one to integrate additional feedback loops or interactions into the development macro view of the system, so that, for example, UX is considered. It also unifies in a natural way agile development and hypothesis-driven development, on the one hand, and iterative, AI-driven algorithms, on the other.

This analysis is however not without cost. There is no tool or standardized approach for formulating or codifying the cybernetic perspective for an item—it requires an ad hoc and not necessarily uniform manual analysis. Also, the different views and components range across first-order, transitional-order, and second-order systems, and identifying important feedback loops and their interactions may be tricky—especially considering that some of these may occur infrequently or as exceptions. Finally, it is also unclear whether the bulk of this analysis can be performed once, or whether or to what extent each new application will require substantial further investigation.

Nonetheless, it is clear that cybernetics affords a coherent perspective on most data science activities and their interactions, from algorithm implementation through to the development cycle and UX.

In future works, other views, some of which are mentioned in this paper, will be investigated as well as a large-scale application to identify feedback loops, interactions, and dependences.

8. REFERENCES

[1] J.R. Laracy and T.J. Marlowe, "A Second-Order Cybernetic Analysis of Hypothesis-Driven Development (HDD)," Keynote presentation, **World Multiconference on Systemics, Cybernetics, and Informatics**, July 20, 2021.

[2] K. Beck, M. Beedle, A. Cockburn, W. Cunningham, M. Fowler, J. Grenning, J. Highsmith, A. Hunt, R. Jeffries, J. Kern, B. Marick, R. C. Martin, S. Mellor, K. Schwaber, J. Sutherland, D. Thomas, and A. van Bennekum, **Manifesto for Agile Software Development,** agilemanifesto.org, 2001.

[3] N. Wiener, **Cybernetics, Second Edition: or the Control and Communication in the Animal and the Machine**, Cambridge, MA: MIT Press, 1965.

[4] C. Shannon, "A Mathematical Theory of Communication," **The Bell System Technical Journal**, Vol. 27, No. 3, July 1948, pp. 379–423.

[5] W.R. Ashby, **An Introduction to Cybernetics**, London: Chapman and Hall, 1956.

[6] J. Von Neumann, A.W. Burks, **Theory of Self-Reproducing Automata**, Urbana: University of Illinois Press, 1966.

[7] H. von Förster and W.R. Ashby, "Biological Computers," **Bioastronautics**, K. E. Schaefer (ed.), New York, Macmillan Co., 1964, pp. 333– 360.

[8] H. von Förster, **Understanding Understanding: Essays on Cybernetics and Cognition**, New York: Springer-Verlag, 2003, p. 289.

[9] R. Abramovitz & H. von Foerster, **Cybernetics of cybernetics: Or, the control and the communication of communication (BCL Report 73.38)**, Biological Computer Laboratory at the University of Illinois, 1974.

[10] M. Mead, "The Cybernetics of Cybernetics," in H. von Foerster, J. D. White, L. J. Peterson, & J. K. Russell (eds.), **Purposive Systems,** Washington, DC: Spartan Books, 1968, pp. 1–11.

[11] G. Pask, B. Scott, & D. Kallikourdis, "A Theory of Conversations and Individuals," **International Journal of Man-Machine Studies**, vol. 5, no. 4, 1973, pp. 443–566.

[12] T. J. Marlowe, J. R. Laracy, J. Fitzpatrick, "Implicit Cybernetic Systems: A Controlling Model, or a Model out of Control?", Keynote presentation, **World Multiconference on Systemics, Cybernetics, and Informatics,** Orlando, FL, USA, July 2021.

[13] J. D. Kelleher and B. Tierney, **Data Science**, Essential Knowledge Series, Cambridge, MA: The MIT Press, 2018.

[14] F. Cady, **The Data Science Handbook,** Hoboken, NJ: Wiley, 2017.

[15] S. S. Skiena, **The Data Science Design Manual,** Cham, Switzerland: Springer, 2017.

[16] M. Kubat, **An Introduction to Machine Learning,** Second Edition, Cham, Switzerland: Springer, 2017.

[17] T. P. Trappenberg, **Fundamentals of Machine Learning,** New York, NY: Oxford University Press, 2020.

[18] J. Han, M. Kamber, J. Pei, **Data Mining: Concepts and Techniques,** Third Edition, Waltham, MA: Morgan Kaufmann, 2012.

[19] T. J. Marlowe, V. Kirova, G. Chang, O. Hashmi, and S. P. Masticola, "Development and Evolution of Agile – Changes in a World of Change", **Journal of Systemics, Cybernetics and Informatics**, Vol. 18, No. 7, 2020.

[20] S. R. Schach, **Object-Oriented and Classical Software Engineering,** Eighth Edition, New York, NY: McGraw Hill, 2010.

[21] H. Cruse, **Neural Networks as Cybernetic Systems,** 2nd and Revised Edition, Bielefeld, Germany: Brains, Minds and Media, 2006.

[22] S. Russell and P. Norvig, **Artificial Intelligence: A Modern Approach,** Third Edition, Upper Saddle River, NJ: Prentice Hall, 2010.

[23] J. D. Kelleher, **Deep Learning,** Essential Knowledge Series, Cambridge, MA: The MIT Press, 2019.

[24] E. Charniak, **Introduction to Deep Learning,** Cambridge, MA: The MIT Press, 2018.

[25] D. E. Holmes, **Big Data: A Very Short Introduction,** New York, NY: Oxford University Press, 2017.

[26] T. Berry, D. Giannakis, and J. Harlim, "Bridging Data Science and Dynamical Systems Theory," **Notices of the American Mathematical Society**, Vol. 67, No. 9, October 2020, pp. 1336–1348.