

# Bridging Disciplinary Divides: Exploring the Synergy of Punctuated Equilibrium Theory and Artificial Neural Networks in Policy Change Analysis

Łukasz Wordliczek

Jagiellonian University in Kraków, Poland

Submitted: July 23, 2023  
Accepted: December 27, 2023

---

## Abstract

*This article explores the conceptual and theoretical intersections between Punctuated Equilibrium Theory (PET) and artificial neural networks (NNs) within the context of policy change analysis. Despite some similarities between PET and NNs, limited systematic research has been conducted to bridge the gap between political science and computer science. The paper addresses this conceptual gap by presenting a theory-oriented, explorative examination, focusing on the commonalities in their principles, such as information processing, dynamic modeling, and adaptation. The study contributes to methodology- and theory-oriented research on policy agendas by extending PET through the incorporation of NNs. The article employs a conceptual lens to establish parallels between PET and NNs, emphasizing their shared features in dealing with complex, dynamic, and adaptive systems. The exploration of anomalies and outliers in policy time-series data serves as a case study to illustrate the potential synergy between political science and STEM sciences (science, technology, engineering, and mathematics). The paper concludes by proposing avenues for future research that can further integrate these allegedly separate disciplines and enhance our understanding of policy dynamics.*

**Keywords:** Punctuated Equilibrium Theory, artificial neural networks, outliers detection, anomaly detection, political science methods, policy change

**DOI:** 10.56583/br.2191

Whoever knows the ways of Nature will more easily notice her deviations; and, on the other hand, whoever knows her deviations will more accurately describe her ways.

Sir Francis Bacon, *Novus Organum*, 1620  
quoted in (Billor, Hadi, and Velleman 2000)

## Introduction

Punctuated Equilibrium Theory (PET) has already secured its place as one of the “classic” approaches in studying policy processes. PET’s theoretical assumptions and empirical findings are insightful for policy students. However, despite its accomplishments, some methodological and theoretical developments are contested (Desmarais 2019; Dowding, Hindmoor, and Martin 2013; Jones 2016; Padgett 1980; Prindle 2006). Therefore, the first objective of the following study is to contribute to the debate on how to study policy changes over time. This allows for a separate study on how punctuations occur in time. For the current purposes, we will deal with quantitatively

---

E-mail addresses and ORCID digital identifiers of the authors

Łukasz Wordliczek • e-mail: lukasz.wordliczek@uj.edu.pl • ORCID: 0000-0003-4043-8587

identifying abrupt changes. The next level of the analysis, which investigates punctuations in substantive detail, is subject to a separate in-depth case study research design.

One of the main contending theories of PET—that is incrementalism—leaving aside its ambiguous meanings (Berry 1990), at the very general level assumes that policy outputs reflect policy inputs regularly and closely to existing levels (Davis, Dempster, and Wildavsky 1974). The state of relative stasis is to a large extent explained by the political culture and complexity of democratic policy-making according to the logic of “muddling through” (Lindblom 1959, 1979). Thus, one could assume that we are between a “dynamic” (PET) and a “static” (incrementalism) approach. But, strictly speaking, PET may be perceived as a corollary to incrementalism, since the former assumes that most of the time policy changes indeed are modest and tend to be “equilibrium seeking” through positive and negative corrections of punctuations (Baumgartner and Jones 2009; Flink and Robinson 2020; Jones and Baumgartner 2005, 112). PET’s added value is that this pattern may, however, change dramatically to cause some severe punctuations. This mixture is dubbed the “dynamic model of choice for public policy” (Jones and Baumgartner 2005).

One of the main features of PET is the description of the policy process in parallel to the information process. In this way, the policy process starts with inputs (sources), and then it goes through throughputs (decisions made according to a bounded rationality scheme and with noise friction), resulting in outputs (irregular and punctuated policy shifts). One of the main reasons for these irregularities is “institutional friction.”

This broad description is similar to generic logic that is present in a quite different field: artificial neural networks (hereafter, the shortened and well-known version is used: neural networks, NNs). Here, one has an input layer (data), hidden layer(s) (processing units) and output layer (results). This layered structure is designed to mimic its biological archetype in the way information/signals are processed. The operation of NNs is based on the presentation of input data, calculation of the values of the model parameters (i.e., “weights”), and setting the output(s). These calculations are possible due to a process called “learning” or “training” of a network, which again resembles the natural world. Out of a gamut of possible learning algorithms, one of the most popular is called “backpropagation.” Simply put, the network is tasked with making the answer (its output) as close as possible to its input. In other words, the main objective is to obtain as small an error as possible. Hence its name, since errors are calculated (i.e., propagated) from the last (back) to the first network’s layer. Such minimizing strategy is performed iteratively through a series of repeated calculations called “epochs.” Interestingly, at the beginning of training a network, random values of weights are presented, and then, through backpropagation, they are fine-tuned to lower the overall error rate. Such a trained model may generalize its performance, i.e. it may use its “knowledge” to challenge new tasks. This feature makes NN flexible, dynamic, and robust—another aspect that resembles the biological original.

Surprisingly, as the similarities between PET and NN may have been acknowledged, very little systematic research has been done on the similarity of their theoretical assumptions. Thus, the paper aims at filling the conceptual gap by bridging the divide between two disciplines: political science and computer science. PET’s and neural networks’ principles make both approaches well-suited for a more detailed examination. Due to such a research perspective, the following piece is going to be theory-oriented, explorative, and introductory, rather than empirical and decisive.

All in all, the study contributes to methodology- and theory-oriented research on policy agendas through some extensions to Punctuated Equilibrium Theory. Since the paper is prepared by a political scientist by training and, presumably, the majority of its prospective readers will be part of the political science community, the substantive argument is supported by as little formal notation as possible. The paper is structured as follows. First, as a gentle introduction to more detailed analysis, basic conceptual similarities between PET and NNs are presented. Second, we proceed with the substantive part of the study by empirical assessment that aims at showing possibilities present in formal modeling of outliers. This serves as a kind of customized case study designed to illustrate possibilities behind merging allegedly different disciplines: political science and STEM sciences. Third, the paper concludes with prospective extensions to the future research agenda.

## 1 Theory of information processes: Punctuated Equilibrium and neural networks

Although there is substantial research on neural networks and Punctuated Equilibrium explanations of policy processes, there is no need to present a detailed literature review. At the same time, however, their common features are almost completely overlooked. There are several reasons not to ignore these similarities.

Let us start with an issue that is worthy of its own detailed studies: biological inspirations. They are manifested in terminology used in analyzing policy processes and neural networks modeling. System, environment, process, signal transmission, noise, equilibrium, stasis, friction, bias, dynamics, feedbacks, adaptation, linearity and nonlinearity—these are just a few of the prominent examples. When we seek deeper, in the Punctuated Equilibrium field, direct links to evolutionary biology are clearly evident (Givel 2010; Jones and Baumgartner 2012; Prindle 2012). When one looks at neural networks, even their own name explicitly evokes their neurophysiological origin. The PET approach is largely based on mechanistic and idiosyncratic explanations of policy formulation and change whereas (artificial) neural networks are devoted to investigation of information processing mechanisms and ways of their formal modeling. Furthermore, biology-inspired features, such as punctuations, anomalies, and novelties, are the main focus here. The last example is a critical factor in the mechanisms of survival and habituation of living organisms, since being able to detect new phenomena means being able to reduce complex information from the environment to the most important question: what is the novelty about? The answer often determines to be a predator or a victim (Marsland 2001, 17). This feature (i.e. the ability to extract information) leads to the next argument.

The second reason for the concurrency of PET and NNs stems from the fact that both approaches emphasize the role of information processing. In both cases, the problem is conceptualized through investigation of collecting, interpreting, prioritizing, and structuring signals from the environment. With different levels of methodological sophistication, the ultimate goal to a large extent stays the same: exploring, extracting and explaining many latent possibilities in the empirically sound data we gather. At the very basic level, to put it bluntly, Punctuated Equilibrium Theory and neural networks are nothing more than examples of distributed information processing systems.<sup>1</sup> Their structural similarities are clearly illustrated in figure 1 (on next page).

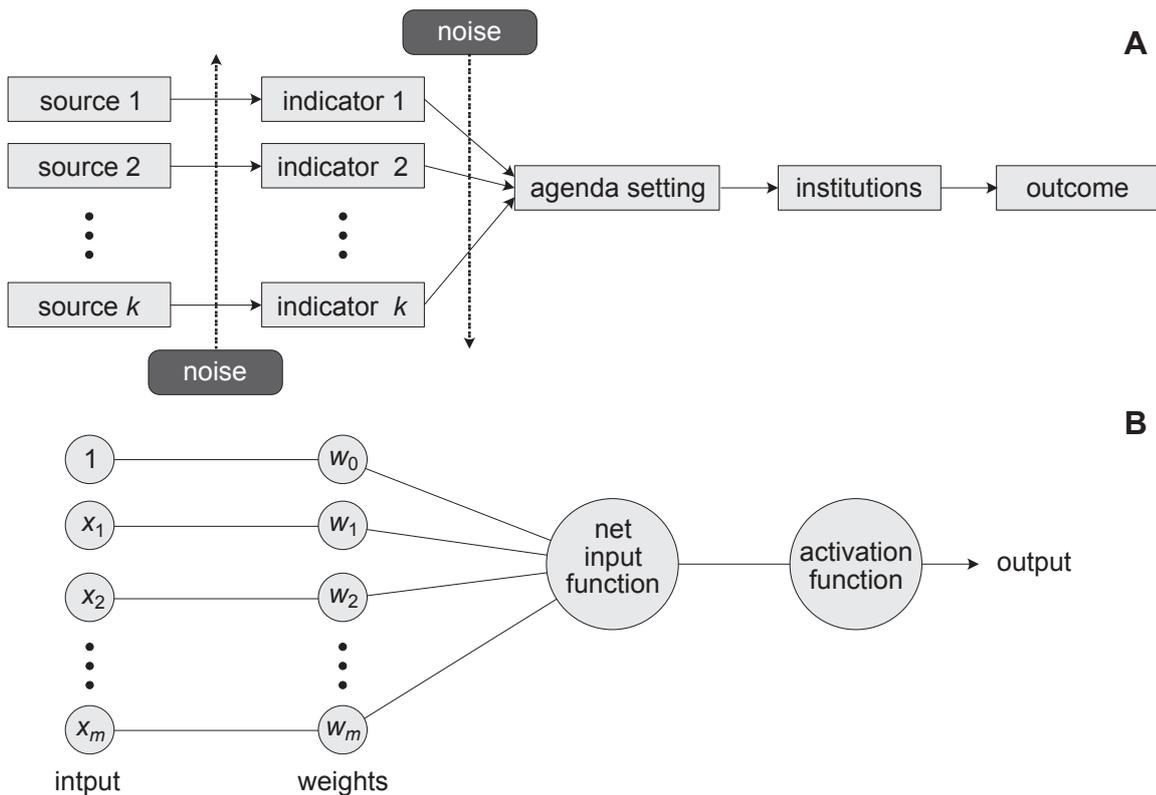
At the same time, this theory is by no means tantamount to trivial research agendas. In the NNs field, for example, the bewildering progress made in a dozen or so years clearly reveals the potential of formal modeling and its practical applications in virtually any aspect of everyday life. Attention paid to the mechanisms of information loss and selective processing is one of the common threads in PET and NNs research. This directly leads to the next point.

Third, both Punctuated Equilibrium Theory and neural networks are indebted to and heavily founded in theoretical assumptions rooted in bounded rationality (Jones 1999, 2003; March and Simon 1958; Padgett 1980; Simon 1955). Policy processes tend to be defined by political actors' selective approach to various issues, combined with their cognitive constraints and biases. Likewise, formal analyses performed by means of ANNs, such as exploratory data analysis, feature engineering, and modeling data, require a careful approach to the limits of data processing tools. This issue is critical at the individual researcher's level, when one is devoted to prioritizing one's own pet ideas over empirical evidence hidden in data at hand. Thus, individual cognitive constraints and biases are indirectly addressed in ANNs. Shortly, it would be hard not to admit that "disruptive patterns of attention" (Hegelich 2017, 66) are present in both approaches.

Fourth, the above point is directly related to a separate argument on many similarities between PET and NNs: both are concerned with systems that are complex (Érdi 2008; Simon 2000). This point relates to the way they conceptualize problems under investigation: efforts aimed at dealing

---

1. Cf. "Information processing involves collecting, assembling, interpreting, and prioritizing signals from the policymaking environment" (Jones and Baumgartner 2012, 7) and "A neural network is a massively parallel distributed processor made up of simple processing units that has a natural propensity for storing experiential knowledge and making it available for use" (Haykin 2009, 2).



**Figure 1.** Information processing systems: policy process (A) and artificial neuron's architecture (B).

Source: Diagram A—(Jones and Baumgartner 2005, 162), diagram B—(Wallace et al. 2017, 120).

with nonlinearities are one of their most evident earmarks. Interestingly, nonlinearity, as a common but unappreciated feature of many systems, was laid as a foundation brick for neural modeling in the pioneering work of researchers as early as in the 1940s (McCulloch and Pitts 1943). To put it succinctly, variables and processes modeled by neural networks do not need to be normally distributed. (In fact, we do not need to have any *ex ante* knowledge of the specific distribution.) By the same token, in many policy processes complex interactions do not trigger a normal distribution of values of interest. Previous research indicates that this is the case of changes in some policy time-series (Jones and Baumgartner 2005). The picture is even more complicated: patterns of government attention also follow “disproportionate information-processing” (Jones and Baumgartner 2005, 5), which makes policy outcomes far from being linear. Complex trajectories are predominantly based on bounded rationality, institutional friction, noise, and limited resources. This leads to a similar but separate argument: Punctuated Equilibrium and many neural networks deal with dynamic systems.

Fifth, consequently, PET and NNs process information signals in a dynamic manner. That is to say that neither of these approaches attributes the same weights to various independent variables but rather models them accordingly. In Punctuated Equilibrium research this argument is expressed as explicitly as it may be: “We think of the traditional political forces such as public opinion, interest groups, elections, and other forms of political participation as providing *weights* for the information signals” (Jones and Baumgartner 2012, 9). Interestingly enough, in the above quotation emphasis is given to the word “weights” by the authors themselves. Also, as we know from the introductory remarks, the processing of neural network signals is tantamount to setting values for the model parameters called weights (according to the ubiquitous convention) (see figure 1). Furthermore, network learning itself is most often a dynamic process (albeit not always).

The next critical feature is that PET and NNs deal with not only complex and dynamic but also adaptive systems. Political systems, just like any social arrangement, respond to exogenous information and are in the process of constant modification. The same is true for neural network

models, which basically calculate their weights: this is nothing more than just changing their values in many iterations to be able to better optimize network performance. To put it succinctly: new information (signal) means a new state (via setting weights). What is not to be ignored is that the result may—and often is—expressed in conditional probability vocabulary.<sup>2</sup>

Seventh, except for other contributions, Punctuated Equilibrium Theory sheds some light on the question of system changes in terms of “underreacting” and “overreacting” to the information delivered. Again, a very similar logic is manifested in neural networks, with the exception that the notions used will be, respectively, “underperformance” and “overfitting.”<sup>3</sup> Therefore, data deficiency and data redundancy are addressed in both approaches.

Last, but not least, some specifics in the approach toward the phenomena under study are also striking. Several arguments have already been mentioned. Here, let us focus on the way information is processed. In Punctuated Equilibrium theory, “Change occurs only when the informational signals from the external world either are extraordinarily strong, on the other hand, or when the signals accumulate over time to overcome the friction. (This latter mechanism is known as error accumulation.) As a consequence, policy-making systems remain stable until the signals from outside exceed a threshold, and then they lurch forward—that is, a policy punctuation occurs; afterward, they resume ‘equilibrium’” (Jones and Baumgartner 2012, 8).

This pattern cannot be more accurate to describe the functioning of biological and, consequently, artificial neurons. Here, neurobiologists and science engineers also talk about “activation threshold,” “potential accumulation,” “signal processing,” and “spikes” (when a neuron fires a signal exceeding, again, some threshold). The relevant point to be added here refers to the fact that both concepts pay attention to parallel processing of inputs. This extension clearly aims to overcome the limitations to serial processing of information typical of many of the more traditional approaches. For a review of the above similarities, please refer to table 1.

**Table 1.** Summary of similarities between Punctuated Equilibrium Theory and artificial neural networks

|   | Punctuated Equilibrium Theory      | Artificial Neural Networks           |
|---|------------------------------------|--------------------------------------|
| Biological inspirations                             | evolution                          | neuroscience                         |
| Main objective                                      | information processing             | information processing               |
| Bounded rationality                                 | directly addressed                 | indirectly addressed                 |
| Characteristic of systems under study               | complexity                         | complexity                           |
| Mode of information processing                      | parallel, dynamic and adaptive     | parallel, dynamic and adaptive       |
| Approach toward data deficiency and data redundancy | “underreacting” and “overreacting” | “underperformance” and “overfitting” |

Thus, it is explicitly argued here that Punctuated Equilibrium and neural networks reveal that both concepts, despite having different origins and separate secured areas in academia, in fact have too much in common to ignore it. Thus, it is even more striking that today there is so little research to bridge the divide that itself may be described as artificial.<sup>4</sup> This leads to the main objective of the following article, which is to contribute to the crossing of the boundaries between two disciplines: political science and computer science. Before considering some justification for the claim, however, let us turn to some necessary clarifications of what PET is and what it is not.

Broadly speaking, the Punctuated Equilibrium is a general theory of policy change. It expands the “standard model” focused on elections, which has a strong normative, theoretical and empirical standing. There are, however, some concerns raised: changes that occur between the elections and/or when there is no electoral change are predominant. Furthermore, the issue of prioritization

2. For Punctuated Equilibrium account see (Jones and Baumgartner 2012, 16). Details on probability-driven explanation in neural networks may be found in any of elementary sources. On this point, as well as many other specifics, see (Garson 1998; Hastie, Tibshirani, and Friedman 2009).

3. The issue is extensively covered in literature. For some basic introductory remarks see, for example, (Colaresi and Mahmood 2017, 198–199).

4. For one of notable examples see (Hegelich 2017).

and framing policy issues is to a large extent underexposed, since the electoral explanation delves mostly into procedural questions. The very same is true of lack of ideological coherence: the “elections matter” argument does not explain why it may be the case that left- or right-wing leaders sometimes pursue goals that are not on their core agenda (Jones and Baumgartner 2012, 5–6). Thus, the above limitations opened the way to searching for other forces that validate policy shifts. One of the possibilities is to reach for the flows of information in the policy trajectory. Interestingly enough, shifts in attention are claimed to be “necessary but not sufficient to bring about policy punctuation” (Jones 2016, 44). Empirical evidence is based mainly on budgetary analysis.

Critical evaluation of Punctuated Equilibrium Theory results in a distinction between two aspects of the policy agenda: *attention* to policy issues and the *content* (i.e., subject) of policy issues (Dowding, Hindmoor, and Martin 2013, 83). Attention is simply a measure of the time that policy actors spend on a given problem. The concept is operationalized in both the Comparative Agendas Project and Policy Agendas Project by such variables as the number of acts passed by the legislature, executive branch position taking, public opinion polls, news media coverage, the US Supreme Court cases, political parties’ manifestos (platforms), etc. The content, on the other hand, seems to be self-evident in light of its name. It is equivalent to specifics included in a given piece of legislation. Thus, attention allocated to policy is not the same as its substantive content. With this argument in mind, one may acknowledge that this has profound consequences for studying policy changes. Indeed, it may be the case that higher levels of attention result in incremental changes and, on the other hand, stasis in time allocation does *not* make the system immune to radical policy punctuations (Dowding, Hindmoor, and Martin 2013, 83).

Any study of the complex trajectories of policy dynamics is determined by a closer investigation of the phenomenon that has already been mentioned: policy change. Thus, before seeking any “explanation of the mechanisms which keep equilibria stable and of those forces or processes which undermine that stability” (Howlett 2009, 246), a research agenda needs to embrace the very concept of data variability, which is the focus of the current article. Before delving into more details, it is necessary, however, to make some reference to scholarship on studying policy change through punctuations.

## 2 Literature review and conceptual appraisal: outliers

The major issue addressed here, broadly speaking, relates to shifts in policy over time. The problem must, however, be conceptualized more rigidly and accurately (Howlett 2009, 243). Having acknowledged the already mentioned discussion on the two aspects of the policy agenda (attention vs. content), it can be argued that the two dimensions are indeed two sides of the same coin *unless* we strive to pinpoint the basic term: policy shift (policy change, policy punctuation, etc.). To put it in other words, it is claimed that policy attention *and* policy content are necessary and sufficient factors in policy changes but *only if* we accordingly define the changes themselves. After all, “incremental changes,” “radical punctuations,” “policy shifts” and “large-scale departures from the past” (Baumgartner, Jones, and Mortensen 2017) seem to be rather taken for granted in previous research, since benchmark sets are as controversial and debatable as they might be. Much of scholarship builds on studying punctuations present in budget data (John and Margetts 2003). Within this research strain, some authors suggested that “punctuations” or “non-incremental changes” in observations are those data points that change more than 40% (Sebók and Berki 2017), more than 30% (Bailey and O’Connor 1975), more or less than 10% (Kemp 1982), above a 20% increase and below a 15% decrease in budget distribution (Jones, Baumgartner, and True 1998), or a 35% increase and a 25% decrease (Jordan 2003). Others established a 5% cut-off threshold of the bottom and top changes of the budget data to be “punctuations” (Baumgartner and Epp 2013). Still others, rather than reach for any user-defined cut points for marking the nonincremental changes, try to superimpose a normal distribution on the empirical distribution of budgetary changes. This approach results in setting thresholds  $< -33\%$  and  $> 35.5\%$  with a margin of  $\pm 5\%$  to indicate, respectively, a negative and a positive punctuation (Flink 2017; Flink and Robinson 2020; Robinson et al. 2007; Robinson, Flink, and King 2013). Obviously, there are at least two limitation

to such a method. First, it presupposes using normally distributed data—an argument that will be critically discussed in some detail below. Second, the above threshold-seeking approaches aim at finding a “global” cutoff across all budget functions, whereas it seems to be more justifiable to differentiate between categories due to their different extreme values instead of setting a generic threshold (Dezhbakhsh, Tohamy, and Aranson 2003, 539–540; Jordan 2003, 352; Munir, Siddiqui, Chattha, et al. 2019, 1997).

Some improvement was suggested in Peter John and Shaun Bevan’s article on punctuations in the legislative agenda of the UK Parliament (John and Bevan 2012). Relying on empirical assessment, the authors proposed a three-category typology of punctuations: procedural, low-salience, and high-salience. Methodologically, instead of paying closer attention to the distributional side of the issue, they applied a simple punctuation measure based on combining percent changes in attention to policy issues across time with the total number of acts passed. However appealing, the limitations of this approach are already known: they are related to identifying punctuations with a user-defined cutoff. John and Bevan set it at the level of 200% yearly changes.

Thus, we are ready to scrutinize the critical problem of large-scale policy changes measurement and, specifically, to address the issue of inappropriateness of setting *any* threshold for punctuations identification by using the percent approach. The necessary preliminary step is to set the relevant context—i.e., describe analytical approaches applicable to variable change research.

Research on meaningful changes in policy has its own, rich tradition (Breunig and Jones 2011; Breunig and Koski 2006; Davis, Dempster, and Wildavsky 1974; Flink and Robinson 2020; Hegelich 2016; Hegelich, Fraune, and Knollmann 2015; Jones et al. 2009; Jones, Baumgartner, and True 1998; Jordan 2003; Padgett 1980; Wildavsky 1964). At the same time, however, up to now relatively little attention has been paid to formally pinpointing the idea of these “meaningful changes”; research has focused mainly on theoretical assumptions and conceptualizing changes in terms of general policy-making processes. This omission is even more glaring if one recalls that significant change in many formal approaches is well defined. Furthermore, the classical—i.e., distributional, approach is highly dependent on the normally distributed variable of interest, whereas, as we already know, this is not always the case.

The above ambiguities are even more important if we move on to conceptual aspects. Examples of the most common terms that are relevant here in describing meaningful variable changes include outliers, anomalies, and novelties.<sup>5</sup> The Merriam-Webster Dictionary<sup>6</sup> defines anomaly as “something different, abnormal, peculiar, or not easily classified . . . deviation from the common rule.” In the subsequent discussion, all these meanings will be addressed. Obviously, when considering semantics, the terms *outlier* (“a statistical observation that is markedly different in value from the others of the sample”) and *novelty* (“something new or unusual”) also apply here.<sup>7</sup> Thus, when looking for any broad definition, one may come up with the following one: an anomaly/outlier is an observation that differs from the rest. This, however, opens questions related to every single word used:

- What is an observation? Is it a point or a series? Is it local or global?
- What does it mean to differ? What measure of deviation should be introduced? Is outlyingness a discrete, binary, 0-1 phenomenon? Or, quite the opposite, is it continuous and can it be assigned a measure?
- What is the rest of the dataset? Where must other observations lie in order not to be considered outliers?

5. Even the very term variable indicates some fluctuations in a given phenomenon. Other terms used in different domains are *change point*, *discordant observation*, *exception*, *aberration*, *surprise*, *peculiarity*, *discord*, *contaminant*, *abnormal object*, *off-nominal operation*, and *atypical observation*. Consequently, the term outlier is often “used by various researchers by different names in different contexts” (Khan and Madden 2014, 347). The original reference relates to one-class classification techniques but the observation is also highly relevant to outliers.

6. See: All definitions mentioned below as published in on-line version of “Merriam Webster Dictionary,” available at <https://www.merriam-webster.com/>.

7. Interestingly enough, the dictionary definition is almost equivalent to the well-known wording proposed in the late 1960s: “An outlying observation, or ‘outlier,’ is one that appears to deviate markedly from other members of the sample in which it occurs” (Grubbs 1969, 1).

The above issues will be addressed below in some detail. As for now, however, one more issue must be clarified. For the current purposes, one should acknowledge that the terms *anomaly* and *outlier* will be used interchangeably from now on to denote observations that differ from others in terms of their features and are relatively rare in data sets in terms of their compliance with statistical modeling. At the same time, however, it is crucial to keep in mind that not every novelty is always an outlier/anomaly, since it may easily fit into a normal observation set. The reason to differentiate between the three terms stems mainly from academic jargon: outliers are usually used in descriptive statistics, whereas anomalies and novelties are associated with other methods. Also, the application domain matters if one speaks of novelties, outliers, or anomalies (Pimentel et al. 2014, 217). Most importantly, novelty detection implies that new observations are accommodated to the already specified data distribution and that they were not included in the training set; in anomaly detection this assumption does not always hold. Bearing this in mind, there is a good reason to use the terms *outlier* and *anomaly* interchangeably in the current analysis (Aggarwal 2017; Munir, Siddiqui, Dengel, et al. 2019).

Consequently, in the subsequent discussion, anomaly detection amounts to “the problem of finding patterns in data that do not conform to expected behavior” (Chandola, Banerjee, and Kumar 2009, 1). It is worthwhile to notice that the definition *does not* refer to any particular variable’s distribution, be it probable or empirical. The reason for such precaution is that intuitively—and sometimes analytically—we are prone to come up with a normal distribution.

And that is the first challenge with anomaly detection: it is a kind of gold standard in descriptive and inferential statistics that since we are dealing with anomalies (or *abnormal* observations), one is particularly apt to think in terms of normally distributed data. However, as will be discussed below, this *normal* logic does not always hold. It has profound analytical consequences. One of the most important is that approaches based on statistical probability distribution are not capable of paying enough attention to the tail of data, the very place where relevant cases (i.e., anomalies) are usually located (Olsson and Holst 2015, 434).

Furthermore, apart from being different from the known behavior and its “theoretical distribution” (Kohonen 2001, 390), anomalies are also relatively rare compared to other observations; if outliers were not that rare, they would be part of the noise or be indistinguishable from normal observations. This makes them even more challenging to study, since the two categories are quite often substantially uneven, which hinders many classical statistical approaches. For this single reason, it may be argued that the whole idea of splitting data into validation and test sets is, at best, precarious, since it imperils anomalies not being included in any of the sets (Maciąg et al. 2021, 10).

There is also the third issue of concern: the way anomalies are identified in terms of data topology. No matter whether one is able to visualize data two-dimensionally or deals with a multidimensional phenomenon, there is a critical need to determine if any observation that is off the modeled data points is indeed an anomaly or just a false alarm (Brotherton and Johnson 2001). Identification of an anomaly is awkward and burdensome, since in many applications we want to know abnormal patterns *before* they emerge. But how to know what anomaly is if it has not been seen before? Logically, it is nothing different from the US Supreme Court’s Justice Steward colloquialism “I know it when I see it.” This leads to the next obstacle: measuring anomalies.

To add the fourth point to the mixture of difficulties in anomaly detection research, one needs to refer to the issue of evaluating the quality of anomaly detectors. Distance- and density-based approaches built on threshold heuristics (see table 2 on page 206) are usually adopted, but soft computing-inspired extensions are also suggested (Gogoi et al. 2011). Also, the well-known confusion matrix may be applied to identify if a model does not find spurious patterns but detects true ones—i.e., a trade-off between, respectively, false positives and true positives is acceptable (Khan and Madden 2014, 347; Tax 2001, 15–16). Applying the right measure is not a trivial question (Hadi, Imon, and Werner 2009; Kriegel, Schubert, and Zimek 2017; Schubert et al. 2012). Consequently, some researchers have suggested building *ensembles* of individual metrics (Aggarwal and Sathe 2017; Benkabou, Benabdeslem, and Canitia 2018; Zimek, Campello, and Sander 2014).

Anomalies/outliers come in different flavors. *Point anomalies* are particular data instances that differ *substantially* from the rest of the data points. On the other hand, *collective anomalies*

are built up of some stacked anomalous data points: for instance, five out of 80 passengers of one of the planes used in 9/11 attacks had features that taken individually would probably have been ignored (e.g., not being US citizens, buying one-way tickets, paying in cash, etc.) but when aggregated they made those passengers substantially different from the other 75 passengers.<sup>8</sup> Also, other empirical analyses showed that, for example, budgetary punctuations may occur in groupings within a window of time (Flink and Robinson 2020; Robinson, Flink, and King 2013). Both point and collective anomalies<sup>9</sup> may also form another category, since cases that have features similar to other observations but differ in terms of a given context in which they appear are accordingly referred to as *contextual anomalies*. Here, a given pattern may be considered normal in a specific process but the very same pattern present in another process may be a manifestation of abnormal behavior. An illustrative and anecdotal example of contextual anomaly is the following statement: “The novelty of the wife in the best friend’s bed lies neither in the wife, nor the friend, nor the bed, but in the unfamiliar conjunction of the three” (O’Keefe and Nadel 1978, 241).

With some introductory intricacies discussed, probably some of the most obvious questions for now would be the following: Why study anomalies at all? Is there any singularity here? Are anomalies idiosyncratic? Are outliers more interesting than inliers? (Hawkins et al. 2002, 170; Williams et al. 2002, 1). Is it not the case that outliers are like fractals, as deleting one of them results in a smaller dataset with its own new outliers? (Abbott 2001, 167–168). While questions like these have rather simple answers in many other domains (security applications, fraud detection, disease diagnosis, engineering, etc.), the issue requires some scrutiny in policy research. Here, it is argued that identifying anomalies is the first and necessary step to find answers to questions that are crucial and more subtle. If we know when abnormal data occur, then we are prone to try to find out why they are present at a given time point. The knowledge gained could be utilized to make the analysis even more reliable by investigating when and why the system departs from its “normality.”

Furthermore, detecting anomalies is highly relevant methodologically. When labeling anomaly observations, one has to decide on the issue of dealing with them. Probably the most obvious and tempting suggestion would be to ignore and eliminate any abnormal data. This used to be a standard procedure for many years due to the sensitivity of many classical approaches to “maverick values” with standard regression at the forefront (Jordan 2003, 351).<sup>10</sup> But is this reject-correct approach accurate (Grubbs 1969, 2)? Is it not too “benign” to treat anomalies as noise? By doing this, which is known as “data cleaning/data cleansing,” one is vulnerable to, albeit not always concerned with, *information cleaning*. Therefore, it is important to consider if these atypical observations do not indicate some meaningful and novel patterns in data. The community concerned with “knowledge discovery in databases” addresses the issue from different perspectives (Breunig et al. 2000; Goldstein and Uchida 2016, 2; Gumbel 1958; Otey, Ghoting, and Parthasarathy 2006, 203). This, almost directly, is expressed in one of the classic definitions of an outlier as “an observation which deviates so much from other observations as to arouse suspicions that it was generated by a different mechanism” (Hawkins 1980, 1). All in all, it might well be the case that we are prone to adjust a phenomenon under study, as well as data, to the methods used. But should it not be the other way around? Should methods and techniques not be determined by data at hand and conceptualized problems? Thus, anomaly detection is perceived here as the very beginning of a more conceptually ambitious and large-scale research agenda on policy dynamics. But how to make the first necessary step?

Recent developments in data science meet the growing needs of researchers. Currently, there are myriad approaches available at different levels of sophistication. Whereas most of them are based in the STEM field, some are derived from such remote areas as immunology. One of the examples

8. See: “Data-Based Detection of Potential Terrorist Attacks on Airplanes.” by Karen Kafadar and Max D. Morris, published on-line by American Statistical Association, section on Statistics in Defense and National Security (<https://community.amstat.org/sdns/home>) [currently not available—Ed.].

9. Point and collective anomalies are not to be confused with local and global anomalies. A local anomaly differs from its local neighbors only, whereas a global anomaly deviates across the entire dataset.

10. This “perennial” neglect of outliers goes back to at least the beginning of the 19th century and Legendre’s work on least squares (Hadi, Imon, and Werner 2009, 57).

suggests looking for the mechanisms of the immune system that differentiate between “self” and “other” in order to apply them to discriminate inliers from novelties through a “negative-selection algorithm” (Dasgupta and Forrest 1996; Forrest et al. 1994; Surace and Worden 2010; Wong, Poll, and KrishnaKumar 2005). Yet, since extensive review of up-to-date research on the subject is beyond the scope of the current analysis and is undertaken elsewhere in detail (Bartkowiak 2011; Chandola, Banerjee, and Kumar 2009; Gogoi et al. 2011; Goldstein and Uchida 2016; Gupta et al. 2014; Hodge and Austin 2004; Khan and Madden 2014; Munir, Siddiqui, Dengel, et al. 2019; Pimentel et al. 2014; Xu, Liu, and Yao 2019), a rather rudimentary survey is presented below.

### 3 Conceptualizing outlier detection

There are at least three research approaches that are relevant to the typology of anomaly detection. The review presented below is partially related to Type 1, Type 2, and Type 3 typologies that are heavily based on the assumption of *ex ante* knowledge on the specific distribution of variables and available training data (Hodge and Austin 2004).<sup>11</sup> Consequently, Type 1 assumes no prior knowledge of data (unsupervised clustering) and usually applies metrics based on densities and distances to evaluate anomaly, Type 2 indicates modeling both normality and abnormality (supervised classification), and Type 3 models only normality or sparsely located abnormality, in which case new data is confronted with the learned model (semi-supervised recognition). For the sake of clarity, the above types may also be clustered into two broader and widely recognized categories: parametric and non-parametric techniques.<sup>12</sup> The issues of anomaly research categorization may also be approached through specific methods used. The following paragraphs concisely merge the three perspectives—i.e., (1) the level of supervision, (2) non/parametric approaches, and (3) the methods used. The rationale is that in many real-life applications, they often intersect.

First, statistical approaches offer some possible measures for outliers (Barnett and Lewis 1994; Grubbs 1969; James et al. 2013, 96–97; Yamanishi et al. 2004). The general idea is to cope with variable densities or probability distributions either in a parametric or nonparametric way. One of the most common parametric approaches is calculating the interquartile range (IQR) (Tukey 1977). It simply means measuring the distance between the first (Q1) and the third (Q3) quartile of a given variable ( $x$ ). An outlier is identified when the following relation is met:

$$(1) \quad x < Q1 - 1.5 \cdot IQR \quad \text{or} \quad Q3 + 1.5 \cdot IQR < x$$

The outlier definition based on the  $1.5 \cdot IQR$  approach is also often illustrated with boxplots (Laurikkala, Juhola, and Kentala 2000). The main issue with the IQR-based measure is well-known: the data must follow a Gaussian distribution, since the logic of the  $1.5 \cdot IQR$  threshold (indicating 99.3% of data points) follows the boundary of  $[\mu - 3\sigma, \mu + 3\sigma]$  technique for Gaussian data that contains 99.7% of observations (Chandola, Banerjee, and Kumar 2009, 30).<sup>13</sup> The idea of treating the data as normally distributed is very tempting, since it is justified by intuition: if we are to look for abnormal data, let us determine a pattern based on statistical properties of a given statistical distribution, and any point that does not fit this pattern is labeled as an anomaly. Unfortunately, in practice the issue is not that straightforward. What is a *normal* pattern? How to define its boundaries? Are they clear-cut or rather fuzzy?<sup>14</sup> Is normality time-invariant and context-invariant? How to deal with noise in data and interdependencies of features? Can we always label normal and abnormal observations? The answers to these questions are too often negative for the intuitive

11. For other typologies of outlier detection see—e.g., (Agyemang, Barker, and Alhaji 2006; Chandola, Banerjee, and Kumar 2009; Goldstein and Uchida 2016; Khan and Madden 2014; Xu, Liu, and Yao 2019).

12. There are, however, interesting attempts at bridging the two domains (Knorr, Ng, and Zamar 2001; Williams et al. 2002). Formally, there is also a semi-parametric technique.

13. Strictly speaking, statistical techniques assume that data is generated from a particular distribution, not necessarily a normal one. This, however, does not help much, since dealing with real data often means noisy and multidimensional data with complex interactions between variables and/or observations.

14. Compare with the observation on the “fuzzy nature of outlyingness” (Hawkins et al. 2002, 171).

and appealing approach to be accepted.<sup>15</sup> (After all, if the issue were that simple, it would not have resulted in such a large number of academic reports, articles, and books.)

Thus, a more nuanced approach is suggested: stochastic models. They are designed to study the whole distribution of policy outcomes. To put it in other words, “Stochastic process models try to ascertain what kinds of probability distributions could have accounted for an observed frequency distribution of outcomes . . . A major reason is the recognition of the key importance of extreme values in the determination of the time paths of complex systems . . . The more traditional regression studies often relegate these critical values to the error structure of ‘unexplained variance’” (Jones, Sulkin, and Larsen 2003, 152).

While being informative in broad-frame studies, the stochastic approach has at least one major limitation: it is not able to explain why specific punctuations occur in a given year. Therefore, the application of this method in investigating attention shifts and agenda setting is questioned and some remedies are suggested (Hegelich, Fraune, and Knollmann 2015). All in all, “policy punctuations that are analyzed as part of the distribution rather than in context” more often than not hinder our “understanding [of] the nature of transition points and the underlying factors associated with large changes in policy” (John and Bevan 2012, 90).

The main problem with many approaches based on supervised learning is that they match data with already known patterns. Through labeling examples in a training data set as being either anomalies or not, the researcher is able to find them relatively easily. Therefore, for example, backpropagation neural networks have been exploited for such problems (Augusteijn and Folkert 2002; Oliveira, Neto, and Meira 2004; Vasconcelos 1995). This practice becomes a critical issue, however, when a new (i.e., unknown, unexpected, and unlabeled) pattern emerges. For supervised frameworks such data will be—for better or worse—just put into already available categories or ignored as a noise (the “data cleaning/data cleansing” dilemma). This lack of complete knowledge of novel instances calls many supervised techniques into question (Brotherton and Johnson 2001; Ma and Perkins 2003, 613) since, rather than try to find a place for a test observation in already predefined categories, one may want “to decide if it belongs to a particular class or not” (Khan and Madden 2014, 345). In such a scenario, another strain of research may be helpful: a family of nonparametric and unsupervised approaches. They may be relevant for studying variable anomalies through not incorporating too many assumptions (Hawkins et al. 2002, 178). The argument seems to be well-grounded, since previous empirical research indicates that, for example, budget changes do not follow a normal distribution but are heavy-tailed distributed (Breunig and Jones 2011; Jones et al. 2009). Thus, as already mentioned, lack of normality makes some techniques a more viable option, since they do not need to provide any knowledge on the distribution of a variable. Furthermore, neural networks are able to “deal with the occurrence of patterns that do not share real membership with any of the training classes” (Vasconcelos 1995, 25), may “automatically discover complex features without having any domain knowledge” (Munir, Siddiqui, Dengel, et al. 2019), and are “capable of learning complex class boundaries” (Hodge and Austin 2004, 108). These attributes make NNs well-suited to solving the distribution-undetermined problem. The selected applications are covered in table 2 (on next page), which sheds some light on differences and similarities in available approaches.<sup>16</sup> Since the paper is not aimed at comprehensive review of these methods, the table serves illustrative purposes only.

---

15. Some authors suggest setting certain thresholds (+/– 10% of a given original data point) to pinpoint anomalies when analyzing time series with more sophisticated approaches (Oliveira, Neto, and Meira 2004, 845–846). Others search for automated ways of finding thresholds based on parametric and non-parametric logic. For the former, deviation is accommodated, whereas for the latter, density distribution may be used (Laptev, Amizadeh, and Flint 2015). Notwithstanding the solution, setting the right threshold is critical for model performance (Xu, Liu, and Yao 2019, 5).

16. A comprehensive survey may be found for example in (Marsland 2003; Pimentel et al. 2014). For attempts towards combining statistical and neural modeling see for example (Ho, Xie, and Goh 2002; Munir, Siddiqui, Chattha, et al. 2019). For the last several years, the hierarchical temporal memory (HTM) approach has also been suggested for modeling temporal sequences (Cui, Ahmad, and Hawkins 2016) and “HTM Applications.” <https://numenta.com/machine-intelligence-technology/applications/> (accessed 2020-01-24). Its performance is, however, arguable (Struye and Latré 2020).

**Table 2.** Selected approaches to anomaly/novelty detection with neural networks

| Author(s), year                              | NN architecture   | “Outlyingness” measure   | Learning formula |
|--|---|--|------------------|
| Sykacek, 1997                                | MLP   | residuals outside classifier’s error bars  | supervised       |
| Brotherton and Johnson, 2001                 | RBF   | Mahalanobis distance   | unsupervised     |
| Hawkins et al., 2002                         | replicator NN (autoencoder)   | mean reconstruction error (mean square error)  | semi-supervised  |
| Williams et al., 2002                        | replicator NN (autoencoder)   | mean reconstruction error (mean square error)  | semi-supervised  |
| Augusteijn and Folkert, 2002                 | probabilistic NN  | probability-based  | supervised       |
| Augusteijn and Folkert, 2002                 | MLP   | Euclidean distance   | supervised       |
| Oliveira, Neto, and Meira, 2004              | MLP + RBF   | area around time series windows  | supervised       |
| Oliveira, and Meira, 2006                    | MLP + recurrent (Elman)   | robust confidence intervals  | supervised       |
| Barreto and Aguayo, 2009                     | competitive neural networks   | decision thresholds based on quantization errors   | unsupervised     |
| Lavin and Ahmad, 2015                        | HTM   | threshold heuristics: scoring function of anomaly windows                                | unsupervised     |
| Marchi et al., 2015                          | LSTM  | reconstruction error   | semi-supervised  |
| Malhorta et al. 2015                         | stacked LSTM  | prediction error   | semi-supervised  |
| Bontemps et al., 2016                        | LSTM  | relative error threshold   | semi-supervised  |
| Kanarachos et al., 2017                      | deep temporal NN + discrete wavelet analysis + Hilbert transformation | Receiver Operating Characteristics (ROC)   | semi-supervised  |
| Hasani, 2017                                 | HTM   | threshold heuristics: probabilistic model of prediction error → likelihood of an anomaly | unsupervised     |
| Ahmad et al., 2017                           | HTM   | threshold heuristics: probabilistic model of prediction error → likelihood of an anomaly | unsupervised     |
| Wu, Zeng, and Yan, 2018                      | HTM   | threshold heuristics: probabilistic model of prediction error → likelihood of an anomaly | unsupervised     |
| Rodriguez, Kotagiri, and Buyya, 2018         | HTM   | threshold heuristics: probabilistic model of prediction error → likelihood of an anomaly | unsupervised     |
| Amarbayasgalan, Jargalsaikhan, and Ryu, 2018 | deep learning autoencoder   | reconstruction error + optimal outlier threshold from retrained model                    | unsupervised     |
| Munir, Siddiqui, Dengel, et al. 2019         | convolutional NN  | distance between the predicted value and the actual value based on Euclidean distance    | unsupervised     |
| Munir, Siddiqui, Chattha, et al. 2019        | convolutional NN  | distance between the predicted value and the actual value based on Euclidean distance    | unsupervised     |

*Note:* Entries are presented in a chronological order.

*Abbreviations:* NN—neural network; MLP—multilayer perceptron, RBF—radial basis function, HTM—Hierarchical Temporal Memory, LSTM—Long Short Term Memory.

This flexibility is accompanied by another distinctive feature. Artificial neural networks—as well as other machine learning techniques—are well-suited for a research agenda that is focused not on inference but on accuracy. However interesting, this makes NNs also hard to interpret, since the models often follow the black-box allegory (Hegelich 2016, 117).<sup>17</sup> Notwithstanding the gamut of possible methods based on neural networks, their general logic is based on constructing classifiers through modeling the input data. Then, the reconstruction error is defined as a difference between the test data and the output. This metrics serves as an anomaly score.

## Conclusion

Taking into account the above discussion, it is argued here that each of the anomaly detection techniques has its unique strengths and weaknesses. Thus, the dilemma which one is *the best* may be seen rather as a kind of optimization problem: Is the tool *good enough* (or *better* than its competitors) to be used, considering our data and research design? The argument presented here seems to shed some light on the issue by pointing to the need to inspect the very nature of the process being modeled. It is argued that many policy-oriented data do not contain any arbitrarily predefined “normal” or “abnormal” data points. To put it in other words, it is acknowledged that data points are *somehow* spread across the distribution and some of them just lie *far* from the mean. Obviously, this statement must be operationalized rigorously, since the critical issue is to determine any measure of the “farness.” In order to do so, the pending problem seems to fall in an unsupervised clustering family of techniques.<sup>18</sup> Consequently, the rationale is that we are to come to terms with anomaly detection considering natural (but not necessarily normal) features of data. This would enable us to treat anomalies independently of any arbitrarily set thresholds (Xu, Liu, and Yao 2019, 5). To put it succinctly, in unsupervised networks their nodes compete for the common features present in the input vectors.<sup>19</sup>

Several points mentioned above were intended to show how much a given political science theory (PET) and one of formal modeling techniques (artificial neural networks) have in common in their theoretical appearances—notwithstanding their idiosyncrasies. This seems to pave the way for further studies. Future extensions to the research are needed in several areas, and empirical testing is probably one of the most obvious choices here.

## References

- ABBOTT, A.D. 2001. *Time Matters. On Theory and Method*. Chicago: University of Chicago Press.
- ADADI, A., and M. BERRADA. 2018. “Peeking Inside the Black-Box: a Survey on Explainable Artificial Intelligence (XAI).” *IEEE Access* 6:52138–52160. doi: 10.1109/ACCESS.2018.2870052.
- AGGARWAL, C.C. 2017. “An Introduction to Outlier Analysis.” In *Outlier Analysis*, 1–34. Cham: Springer International Publishing.
- AGGARWAL, C.C., and S. SATHE. 2017. *Outlier Ensembles. An Introduction*. Cham: Springer International Publishing.
- AGYEMANG, M., K. BARKER, and R. ALHAJJ. 2006. “A Comprehensive Survey of Numeric and Symbolic Outlier Mining Techniques.” *Intelligent Data Analysis* 10:521–538. doi: 10.3233/IDA-2006-10604.
- AUGUSTEIJN, M.F., and B.A. FOLKERT. 2002. “Neural Network Classification and Novelty Detection.” *International Journal of Remote Sensing* 23 (14):2891–2902. doi: 10.1080/01431160110055804.

17. This fact notwithstanding, there are some promising attempts at addressing this limitation (Adadi and Berrada 2018; Baró et al. 2018; Ghorbani, Abid, and Zou 2019).

18. Referring back to the Type1 - Type 2 - Type 3 typology, Type 2 approach does not hold here since “classification algorithms require a good spread of both normal and abnormal data—i.e., the data should cover the entire distribution to allow generalisation by the classifier” (Hodge and Austin 2004, 89). Furthermore, Type 3 posits that data must be “pre-classified” and it “requires the full gamut of normality to be available for training to permit generalisation” (Hodge and Austin 2004, 90). As we already know, this is not the case here.

19. Clearly, there are dozens of algorithms used in an unsupervised anomaly detection problems. For an illustrative review see for example (Campos et al. 2016; Goldstein and Uchida 2016).

- BAILEY, J.J., and R.J. O'CONNOR. 1975. "Operationalizing Incrementalism: Measuring the Muddles." *Public Administration Review* 35 (1):60–66. doi: 10.2307/975202.
- BARNETT, V., and T. LEWIS. 1994. *Outliers in Statistical Data*. 3rd ed., Wiley Series in Probability and Mathematical Statistics. Applied Probability and Statistics. Chichester – New York: Wiley.
- BARÓ, X., H.J. ESCALANTE, S. ESCALERA, U. GÜÇLÜ, Y.M. GÜÇLÜTÜRK, I. GUYON, and M. VAN GERVEN, eds. 2018. *Explainable and Interpretable Models in Computer Vision and Machine Learning*. Cham: Springer International Publishing.
- BARTKOWIAK, A. 2011. "Anomaly, Novelty, One-Class Classification: a Comprehensive Introduction." *International Journal of Computer Information Systems and Industrial Management Applications* 3:061–071.
- BAUMGARTNER, F.R., and D.A. EPP. 2013. "Explaining Punctuations." Annual Meetings of the Comparative Agendas Project, Antwerp, Belgium, June 27–29.
- BAUMGARTNER, F.R., and B.D. JONES. 2009. *Agendas and Instability in American Politics*. 2nd ed., Chicago Studies in American Politics. Chicago: The University of Chicago Press.
- BAUMGARTNER, F.R., B.D. JONES, and P.B. MORTENSEN. 2017. "Punctuated Equilibrium Theory: Explaining Stability and Change in Public Policymaking." In *Theories of the Policy Process*, edited by C.M. Weible and P.A. Sabatier, 155–187. New York, NY: Westview Press.
- BENKABOU, S.-E., K. BENABDESLEM, and B. CANITIA. 2018. "Unsupervised Outlier Detection for Time Series by Entropy and Dynamic Time Warping." *Knowledge and Information Systems* 54 (2):463–486. doi: 10.1007/s10115-017-1067-8.
- BERRY, W.D. 1990. "The Confusing Case of Budgetary Incrementalism: Too Many Meanings for a Single Concept." *The Journal of Politics* 52 (1):167–196. doi: 10.2307/2131424.
- BILLOR, N., A.S. HADI, and P.F. VELLEMAN. 2000. "BACON: Blocked Adaptive Computationally Efficient Outlier Nominators." *Computational Statistics & Data Analysis* 34 (3):279–298. doi: 10.1016/S0167-9473(99)00101-2.
- BREUNIG, C., and B.D. JONES. 2011. "Stochastic Process Methods with an Application to Budgetary Data." *Political Analysis* 19 (1):103–117.
- BREUNIG, C., and C. KOSKI. 2006. "Punctuated Equilibria and Budgets in the American States." *Policy Studies Journal* 34 (3):363–379. doi: 10.1111/j.1541-0072.2006.00177.x.
- BREUNIG, M.M., H.-P. KRIEGEL, R.T. NG, and J. SANDER. 2000. "LOF: Identifying Density-Based Local Outliers." Proceedings of the 2000 ACM SIGMOD international conference on Management of data, Dallas, Texas, USA.
- BROTHERTON, T., and T. JOHNSON. 2001. "Anomaly Detection for Advanced Military Aircraft Using Neural Networks." 2001 IEEE Aerospace Conference Proceedings (Cat. No.01TH8542), March 10–17.
- CAMPOS, G.O., A. ZIMEK, J. SANDER, R.J.G.B. CAMPELLO, B. MICENKOVÁ, E. SCHUBERT, I. ASSENT, and M.E. HOULE. 2016. "On the Evaluation of Unsupervised Outlier Detection: Measures, Datasets, and an Empirical Study." *Data Mining and Knowledge Discovery* 30 (4): 891–927. doi: 10.1007/s10618-015-0444-8.
- CHANDOLA, V., A. BANERJEE, and V. KUMAR. 2009. "Anomaly Detection: a Survey." *ACM Computing Surveys* 41 (3):1–58. doi: 10.1145/1541880.1541882.
- COLARESI, M., and Z. MAHMOOD. 2017. "Do the Robot: Lessons from Machine Learning to Improve Conflict Forecasting." *Journal of Peace Research* 54 (2):193–214.
- CUI, Y., S. AHMAD, and J. HAWKINS. 2016. "Continuous Online Sequence Learning with an Unsupervised Neural Network Model." *Neural Computation* 28 (11):2474–2504. doi: 10.1162/NECO\_a\_00893.
- DASGUPTA, D., and S. FORREST. 1996. "Novelty Detection in Time Series Data Using Ideas from Immunology." ISCA 5th International Conference on Intelligent Systems, Reno, USA, June 19–21.
- DAVIS, O.A., M.A.H. DEMPSTER, and A. WILDAVSKY. 1974. "Towards a Predictive Theory of Government Expenditure: Us Domestic Appropriations." *British Journal of Political Science* 4 (4):419–452. doi: 10.1017/S0007123400009650.
- DESMARAIS, B.A. 2019. "Punctuated Equilibrium or Incrementalism in Policymaking: What We Can and Cannot Learn from the Distribution of Policy Changes." *Research & Politics* 6 (3):1–6. doi: 10.1177/2053168019871399.

- DEZHBAKHSH, H., S.M. TOHAMY, and P.H. ARANSON. 2003. "A New Approach for Testing Budgetary Incrementalism." *Journal of Politics* 65 (2):532–558. doi: 10.1111/1468-2508.t01-3-00014.
- DOWDING, K., A. HINDMOOR, and A. MARTIN. 2013. "Australian Public Policy: Attention, Content and Style." *Australian Journal of Public Administration* 72 (2):82–88. doi: 10.1111/1467-8500.12012.
- ÉRDI, P.T. 2008. *Complexity Explained*, Springer Complexity. Berlin: Springer.
- FLINK, C.M. 2017. "Rethinking Punctuated Equilibrium Theory: a Public Administration Approach to Budgetary Changes." *Policy Studies Journal* 45 (1):101–120. doi: 10.1111/psj.12114.
- FLINK, C.M., and S.E. ROBINSON. 2020. "Corrective Policy Reactions: Positive and Negative Budgetary Punctuations." *Journal of Public Policy* 40 (1):96–115. doi: 10.1017/S0143814X18000259.
- FORREST, S., A.S. PERELSON, L. ALLEN, and R. CHERUKURI. 1994. "Self-Nonself Discrimination in a Computer." IEEE Computer Society Symposium on Research in Security and Privacy, Oakland, CA, May 16–18.
- GARSON, G.D. 1998. *Neural Networks. An Introductory Guide for Social Scientists*, New Technologies for Social Research. London – Thousand Oaks, CA: Sage.
- GHORBANI, A., A. ABID, and J. ZOU. 2019. "Interpretation of Neural Networks Is Fragile." *Proceedings of the AAAI Conference on Artificial Intelligence* 33 (01):3681–3688. doi: 10.1609/aaai.v33i01.33013681.
- GIVEL, M. 2010. "The Evolution of the Theoretical Foundations of Punctuated Equilibrium Theory in Public Policy." *Review of Policy Research* 27 (2):187–198. doi: 10.1111/j.1541-1338.2009.00437.x.
- GOGOI, P., D.K. BHATTACHARYYA, B. BORAH, and J.K. KALITA. 2011. "A Survey of Outlier Detection Methods in Network Anomaly Identification." *The Computer Journal* 54 (4):570–588. doi: 10.1093/comjnl/bxr026.
- GOLDSTEIN, M., and S. UCHIDA. 2016. "A Comparative Evaluation of Unsupervised Anomaly Detection Algorithms for Multivariate Data." *PLoS One* 11 (4):1–31 (e0152173). doi: 10.1371/journal.pone.0152173.
- GRUBBS, F.E. 1969. "Procedures for Detecting Outlying Observations in Samples." *Technometrics* 11 (1):1–21. doi: 10.2307/1266761.
- GUMBEL, E.J. 1958. *Statistics of Extremes*. New York: Columbia University Press.
- GUPTA, M., J. GAO, C.C. AGGARWAL, and J. HAN. 2014. "Outlier Detection for Temporal Data: a Survey." *IEEE Transactions on Knowledge and Data Engineering* 26 (9):2250–2267. doi: 10.1109/TKDE.2013.184.
- HADI, A.S., A.H.M.R. IMON, and M. WERNER. 2009. "Detection of Outliers." *WIREs Computational Statistics* 1 (1):57–70. doi: 10.1002/wics.6.
- HALDANE, J.B.S. 1928. *Possible Worlds, and Other Papers*. New York – London: Harper & Brothers.
- HASTIE, T., R. TIBSHIRANI, and J.H. FRIEDMAN. 2009. *The Elements of Statistical Learning. Data Mining, Inference, and Prediction*. 2nd ed., Springer Series in Statistics. New York, NY: Springer.
- HAWKINS, D.M. 1980. *Identification of Outliers*, Monographs on Applied Probability and Statistics. London – New York: Chapman and Hall.
- HAWKINS, S., H. HE, G. WILLIAMS, and R. BAXTER. 2002. "Outlier Detection Using Replicator Neural Networks." In *Data Warehousing and Knowledge Discovery*, edited by Y. Kambayashi, W. Winiwarter and M. Arikawa, 170–180. Berlin – Heidelberg: Springer Berlin Heidelberg.
- HAYKIN, S.S. 2009. *Neural Networks and Learning Machines*. 3rd ed. New York: Prentice Hall.
- HEGELICH, S. 2016. "Decision Trees and Random Forests: Machine Learning Techniques to Classify Rare Events." *European Policy Analysis* 2 (1):98–120. doi: 10.18278/epa.2.1.7.
- HEGELICH, S. 2017. "Deep Learning and Punctuated Equilibrium Theory." *Cognitive Systems Research* 45:59–69. doi: 10.1016/j.cogsys.2017.02.006.
- HEGELICH, S., C. FRAUNE, and D. KNOLLMANN. 2015. "Point Predictions and the Punctuated Equilibrium Theory: a Data Mining Approach—U.S. Nuclear Policy as Proof of Concept." *Policy Studies Journal* 43 (2):228–256. doi: 10.1111/psj.12089.
- HO, S.L., M. XIE, and T.N. GOH. 2002. "A Comparative Study of Neural Network and Box-Jenkins ARIMA Modeling in Time Series Prediction." *Computers & Industrial Engineering* 42 (2): 371–375. doi: 10.1016/S0360-8352(02)00036-0.

- HODGE, V., and J. AUSTIN. 2004. "A Survey of Outlier Detection Methodologies." *Artificial Intelligence Review* 22 (2):85–126. doi: 10.1023/B:AIRE.0000045502.10941.a9.
- HOWLETT, M. 2009. "Process Sequencing Policy Dynamics: beyond Homeostasis and Path Dependency." *Journal of Public Policy* 29 (3):241–262.
- JAMES, G., D. WITTEN, T. HASTIE, and R. TIBSHIRANI. 2013. *An Introduction to Statistical Learning. With Applications in R*, Springer Texts in Statistics. New York: Springer.
- JOHN, P., and S. BEVAN. 2012. "What Are Policy Punctuations? Large Changes in the Legislative Agenda of the UK Government, 1911–2008." *Policy Studies Journal* 40 (1):89–108. doi: 10.1111/j.1541-0072.2011.00435.x.
- JOHN, P., and H. MARGETTS. 2003. "Policy Punctuations in the UK: Fluctuations and Equilibria in Central Government Expenditure Since 1951." *Public Administration* 81 (3):411–432. doi: 10.1111/1467-9299.00354.
- JONES, B.D. 1999. "Bounded Rationality." *Annual Review of Political Science* 2 (1):297–321. doi: 10.1146/annurev.polisci.2.1.297.
- JONES, B.D. 2003. "Bounded Rationality and Political Science: Lessons from Public Administration and Public Policy." *Journal of Public Administration Research and Theory: J-PART* 13 (4):395–412.
- JONES, B.D. 2016. "The Comparative Policy Agendas Projects as Measurement Systems: Response to Dowding, Hindmoor and Martin." *Journal of Public Policy* 36 (1):31–46. doi: 10.1017/S0143814X15000161.
- JONES, B.D., and F.R. BAUMGARTNER. 2005. *The Politics of Attention. How Government Prioritizes Problems*. Chicago: University of Chicago Press.
- JONES, B.D., and F.R. BAUMGARTNER. 2012. "From There to Here: Punctuated Equilibrium to the General Punctuation Thesis to a Theory of Government Information Processing." *Policy Studies Journal* 40 (1):1–20. doi: 10.1111/j.1541-0072.2011.00431.x.
- JONES, B.D., F.R. BAUMGARTNER, C. BREUNIG, C. WLEZIEN, S. SOROKA, M. FOUCAULT, A. FRANÇOIS, C. GREEN-PEDERSEN, C. KOSKI, P. JOHN, P.B. MORTENSEN, F. VARONE, and S. WALGRAVE. 2009. "A General Empirical Law of Public Budgets: a Comparative Analysis." *American Journal of Political Science* 53 (4):855–873.
- JONES, B.D., F.R. BAUMGARTNER, and J.L. TRUE. 1998. "Policy Punctuations: U.S. Budget Authority, 1947–1995." *The Journal of Politics* 60 (1):1–33. doi: 10.2307/2647999.
- JONES, B.D., T. SULKIN, and H.A. LARSEN. 2003. "Policy Punctuations in American Political Institutions." *The American Political Science Review* 97 (1):151–169.
- JORDAN, M.M. 2003. "Punctuations and Agendas: a New Look at Local Government Budget Expenditures." *Journal of Policy Analysis and Management* 22 (3):345–360.
- KEMP, K.A. 1982. "Instability in Budgeting for Federal Regulatory Agencies." *Social Science Quarterly* 63 (4):643–660.
- KHAN, S.S., and M.G. MADDEN. 2014. "One-Class Classification: Taxonomy of Study and Review of Techniques." *The Knowledge Engineering Review* 29 (3):345–374. doi: 10.1017/S026988891300043X.
- KNORR, E.M., R.T. NG, and R.H. ZAMAR. 2001. "Robust Space Transformations for Distance-Based Operations." 7th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, August 26–29.
- KOHONEN, T. 2001. *Self-Organizing Maps*. 3rd ed., Springer Series in Information Sciences. Berlin – New York: Springer.
- KRIEGEL, H.-P., E. SCHUBERT, and A. ZIMEK. 2017. "The (Black) Art of Runtime Evaluation: Are We Comparing Algorithms or Implementations?" *Knowledge and Information Systems* 52 (2):341–378. doi: 10.1007/s10115-016-1004-2.
- LAPTEV, N., S. AMIZADEH, and I. FLINT. 2015. "Generic and Scalable Framework for Automated Time-Series Anomaly Detection." 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Sydney, NSW, Australia, August 10–13.
- LAURIKKALA, J., M. JUHOLA, and E. KENTALA. 2000. "Informal Identification of Outliers in Medical Data." 5th International Workshop on Intelligent Data Analysis in Medicine and Pharmacology [held] at the 14th European Conference on Artificial Intelligence (ECAI-2000), Berlin, Germany, August 20–25.
- LINDBLOM, C.E. 1959. "The Science of 'Muddling Through'." *Public Administration Review* 19 (2):79–88. doi: 10.2307/973677.

- LINDBLOM, C.E. 1979. "Still Muddling, Not Yet Through." *Public Administration Review* 39 (6): 517–526. doi: 10.2307/976178.
- MA, J., and S. PERKINS. 2003. "Online Novelty Detection on Temporal Sequences." 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Washington, D.C., August 24–27.
- MACIĄG, P.S., M. KRYSZKIEWICZ, R. BEMBENIK, J. L. LOBO, and J. DEL SER. 2021. "Unsupervised Anomaly Detection in Stream Data with Online Evolving Spiking Neural Networks." *Neural Networks* 139:118–139. doi: 10.1016/j.neunet.2021.02.017.
- MARCH, J.G., and H.A. SIMON. 1958. *Organizations*. New York: Wiley.
- MARSLAND, S. 2001. "On-Line Novelty Detection through Self-Organisation, with Application to Inspection Robotics." doctoral thesis, Department of Computer Science, University of Manchester.
- MARSLAND, S. 2003. "Novelty Detection in Learning Systems." *Neural Computing Surveys* 3:157–195.
- MCCULLOCH, W.S., and W. PITTS. 1943. "A Logical Calculus of the Ideas Immanent in Nervous Activity." *The Bulletin of Mathematical Biophysics* 5 (4):115–133. doi: 10.1007/BF02478259.
- MUNIR, M., S.A. SIDDIQUI, M.A. CHATTHA, A. DENGEL, and S. AHMED. 2019. "FuseAD: Unsupervised Anomaly Detection in Streaming Sensors Data by Fusing Statistical and Deep Learning Models." *Sensors* 19 (11):2451. doi: 10.3390/s19112451.
- MUNIR, M., S.A. SIDDIQUI, A. DENGEL, and S. AHMED. 2019. "DeepAnT: a Deep Learning Approach for Unsupervised Anomaly Detection in Time Series." *IEEE Access* 7:1991–2005. doi: 10.1109/ACCESS.2018.2886457.
- O'KEEFE, J., and L. NADEL. 1978. *The Hippocampus as a Cognitive Map*. Oxford – New York: Clarendon Press; Oxford University Press.
- OLIVEIRA, A.L.I., F.B.D.L. NETO, and S.R.D.L. MEIRA. 2004. "Combining MLP and RBF Neural Networks for Novelty Detection in Short Time Series." MICAI 2004: Advances in Artificial Intelligence, Third Mexican International Conference on Artificial Intelligence, Mexico City, Mexico, April 26–30.
- OLSSON, T., and A. HOLST. 2015. "A Probabilistic Approach to Aggregating Anomalies for Unsupervised Anomaly Detection with Industrial Applications." 28th International Florida Artificial Intelligence Research Society Conference, FLAIRS 2015, Hollywood, FL, USA, May 18–20.
- OTEY, M.E., A. GHOTING, and S. PARTHASARATHY. 2006. "Fast Distributed Outlier Detection in Mixed-Attribute Data Sets." *Data Mining and Knowledge Discovery* 12 (2):203–228. doi: 10.1007/s10618-005-0014-6.
- PADGETT, J.F. 1980. "Bounded Rationality in Budgetary Research." *The American Political Science Review* 74 (2):354–372. doi: 10.2307/1960632.
- PIMENTEL, M.A.F., D.A. CLIFTON, L. CLIFTON, and L. TARASSENKO. 2014. "A Review of Novelty Detection." *Signal Processing* 99:215–249. doi: 10.1016/j.sigpro.2013.12.026.
- PRINDLE, D. 2006. "Stephen Jay Gould as a Political Theorist." *Politics and the Life Sciences* 25 (1/2):2–14.
- PRINDLE, D.F. 2012. "Importing Concepts from Biology into Political Science: The Case of Punctuated Equilibrium." *Policy Studies Journal* 40 (1):21–44. doi: 10.1111/j.1541-0072.2011.00432.x.
- ROBINSON, S., E., C. FLOUN'SAY, K.J. MEIER, and L.J. O'TOOLE JR. 2007. "Explaining Policy Punctuations: Bureaucratization and Budget Change." *American Journal of Political Science* 51 (1):140–150.
- ROBINSON, S.E., C.M. FLINK, and C.M. KING. 2013. "Organizational History and Budgetary Punctuation." *Journal of Public Administration Research and Theory* 24 (2):459–471. doi: 10.1093/jopart/mut035.
- SCHUBERT, E., R. WOJDANOWSKI, A. ZIMEK, and H.-P. KRIEGEL. 2012. "On Evaluation of Outlier Rankings and Outlier Scores." 12th SIAM International Conference on Data Mining (SDM), Anaheim, CA, USA, Apr 26–28.
- SEBÓK, M., and T. BERKI. 2017. "Incrementalism and Punctuated Equilibrium in Hungarian Budgeting (1991–2013)." *Journal of Public Budgeting, Accounting & Financial Management* 29 (2):151–180. doi: 10.1108/JPBAFM-29-02-2017-B001.
- SIMON, H. 2000. "Public Administration in Today's World of Organizations and Markets." *PS: Political Science & Politics* 33 (4):749–756. doi: 10.2307/420911.
- SIMON, H.A. 1955. "A Behavioral Model of Rational Choice." *The Quarterly Journal of Economics* 69 (1):99–118. doi: 10.2307/1884852.

- STRUYE, J., and S. LATRÉ. 2020. “Hierarchical Temporal Memory and Recurrent Neural Networks for Time Series Prediction: an Empirical Validation and Reduction to Multilayer Perceptrons.” *Neurocomputing* 396:291–301. doi: 10.1016/j.neucom.2018.09.098.
- SURACE, C., and K. WORDEN. 2010. “Novelty Detection in a Changing Environment: a Negative Selection Approach.” *Mechanical Systems and Signal Processing* 24 (4):1114–1128. doi: 10.1016/j.ymssp.2009.09.009.
- TAX, D.M.J. 2001. “One-Class Classification: Concept Learning in the Absence of Counter-Examples.” doctoral thesis, Technische Universiteit Delft (The Netherlands).
- TUKEY, J.W. 1977. *Exploratory Data Analysis*, Addison-Wesley Series in Behavioral Science. Reading, Mass.: Addison-Wesley Pub. Co.
- VASCONCELOS, G.C. 1995. “An Investigation of Feedforward Neural Networks with Respect to the Detection of Spurious Patterns.” doctoral thesis, University of Kent at Canterbury.
- WALLACE, B., S. AKHAVAN-MASOULEH, A. DAVIS, M. WOJNOWICZ, and J.H. BROCK (CYLANCE DATA SCIENCE TEAM). 2017. *Introduction to Artificial Intelligence for Security Professionals*. Irvine, CA: The Cylance Press.
- WILDAVSKY, A.B. 1964. *The Politics of the Budgetary Process*. Boston: Little.
- WILLIAMS, G., R. BAXTER, H. HE, S. HAWKINS, and L. GU. 2002. “A Comparative Study of RNN for Outlier Detection in Data Mining.” 2002 IEEE International Conference on Data Mining, Maebashi City, Japan, December 9–12.
- WONG, D., S. POLL, and K. KRISHNAKUMAR. 2005. “Aircraft Fault Detection and Classification Using Multi-Level Immune Learning System.” Infotech@Aerospace Conference, Arlington, Virginia, September 26–29.
- XU, X., H. LIU, and M. YAO. 2019. “Recent Progress of Anomaly Detection.” *Complexity* 2019:2686378. doi: 10.1155/2019/2686378.
- YAMANISHI, K., J.-I. TAKEUCHI, G. WILLIAMS, and P. MILNE. 2004. “On-Line Unsupervised Outlier Detection Using Finite Mixtures with Discounting Learning Algorithms.” *Data Mining and Knowledge Discovery* 8 (3):275–300. doi: 10.1023/B:DAMI.0000023676.72185.7c.
- ZIMEK, A., R.J.G.B. CAMPELLO, and J. SANDER. 2014. “Ensembles for Unsupervised Outlier Detection: Challenges and Research Questions a Position Paper.” *ACM SIGKDD Explorations Newsletter* 15 (1):11–22. doi: 10.1145/2594473.2594476.