

EMPIRIA

EMPIRIA. Revista de Metodología de las Ciencias Sociales

ISSN: 1139-5737

ISSN: 2174-0682

empiria@poli.uned.es

Universidad Nacional de Educación a Distancia

España

Wordliczek, Lukasz

Neural Networks and Political Science: Testing the Methodological Frontiers*

EMPIRIA. Revista de Metodología de las Ciencias Sociales, vol. 57, 2023, Enero-, pp. 37-62

Universidad Nacional de Educación a Distancia

Madrid, España

DOI: <https://doi.org/10.5944/empiria.57.2023.36429>

Disponible en: <https://www.redalyc.org/articulo.oa?id=297176235002>

- ▶ [Cómo citar el artículo](#)
- ▶ [Número completo](#)
- ▶ [Más información del artículo](#)
- ▶ [Página de la revista en redalyc.org](#)

redalyc.org

Sistema de Información Científica Redalyc

Red de Revistas Científicas de América Latina y el Caribe, España y Portugal

Proyecto académico sin fines de lucro, desarrollado bajo la iniciativa de acceso abierto

*Neural Networks and Political Science: Testing the Methodological Frontiers**

Redes neuronales y ciencia política: probando las fronteras metodológicas

ŁUKASZ WORDLICZEK

Jagiellonian University
lukasz.wordliczek@uj.edu.pl (POLONIA)
<https://orcid.org/0000-0003-4043-8587>

Recibido: 14.06. 2021

Aceptado: 09.10.2022

ABSTRACT

In recent years, a number of significant methodological re-evaluations have taken place in various disciplines of science due to machine learning developments. This is particularly evident in STEM disciplines, while the behavioral and social sciences seem to approach these phenomena with some reserve. A good example is the use of artificial neural networks. Yet, acknowledging their characteristics, it can be safely assumed that they are relatively well designed to solve many problems in political science. This is due to the nature of many social phenomena that are characterized by at least three features: (1) their theoretical basis is not ultimately determined, (2) they lack fully recognized functional relations, and (3) they are described by data that occur in a form that may be cumbersome for traditional modeling. Therefore, the article proceeds with some encouragement for the use of neural networks. At the same time, however, we need to proceed with caution. To mitigate possible opacity, a new political science-informed conceptualization of neural networks categorization scheme is proposed. This aims to help social scientists come to terms with one of the exponentially developing methods in the machine learning toolbox.

KEYWORDS

Artificial neural networks, political science methodology, machine learning, prediction, classification.

*This work was supported by the National Science Centre, Poland [grant number UMO-2014/15/B/HS5/01885]. The article originates in the argument initially presented in a book published by the author in Polish (Łukasz Wordliczek, *Wykorzystanie sztucznych sieci neuronowych w badaniach politologicznych*, Warszawa: Wydawnictwo Naukowe PWN, 2020, 252 p.).

RESUMEN

En los últimos años, se han llevado a cabo una serie de reevaluaciones metodológicas significativas en varias disciplinas de la ciencia, debido a los desarrollos del aprendizaje automático. Esto es particularmente evidente en las matemáticas y la informática, mientras que en las ciencias sociales y del comportamiento estos fenómenos parecen abordarse con cierta reserva. Un buen ejemplo es el uso de redes neuronales artificiales. Sin embargo, reconociendo estas características, se puede suponer que están relativamente bien diseñados para resolver muchos problemas de la ciencia política. Esto se debe a la naturaleza de los fenómenos sociales que, al menos, se caracterizan por lo siguiente: (1) Su base teórica no está en última instancia determinada, (2) Carecen de relaciones funcionales plenamente reconocidas y (3) Se describen mediante datos que de manifestarse en cierta forma puede ser engorroso para el modelado tradicional. El presente artículo insiste en los estímulos para el uso de redes neuronales, aunque reconocemos que se debe proceder con cautela. Para mitigar la posible opacidad, se propone una nueva conceptualización del esquema de categorización de redes neuronales basada en la ciencia política. Esto tiene como objetivo ayudar a los científicos sociales a aceptar uno de los métodos de desarrollo exponencial en esa caja de herramientas que es el aprendizaje automático.

PALABRAS CLAVE

Redes neuronales artificiales, metodología de las ciencias políticas, aprendizaje automático, predicción, clasificación.

1. INTRODUCTION

In recent years, we have witnessed at least three phenomena that pose challenges to traditionally practiced science: the spread of computers, the data deluge, and the development of analytical techniques. In STEM disciplines (Science, Technology, Engineering, Mathematics) this state is considered to be quite 'natural', but in the case of the behavioral and social sciences, not to mention the humanities, it often provokes some reservation but also, most importantly, misunderstanding and misgiving. The question is: should this state of affairs be considered satisfactory? Does it really mean that current methodological developments do not encourage going further beyond the traditional boundaries of scientific disciplines? Why are so many political scientists behind the curve? And finally, could it be that the particular analytical solutions that have been developed in the STEM disciplines do not bear due reference to research conducted

in other areas? One of the analytical approaches that can be contextualized here is artificial neural networks.¹


The reader should be warned against treating the following article as a comprehensive exploration of various literatures. It rather puts in context the following case studies that were rigorously chosen to address the above questions with a ‘political science friendly’ categorization scheme. Thus, the aim here is to offer more than just a simple literature review. In fact, the methodological approach covered here serves more as a vehicle to discuss multifaceted challenges in contemporary research. Thus, it is argued that the below discussion would better serve political/social science practitioners, whereas it also may be illustrative for STEM researchers in terms of reconsidering their well-acknowledged gold standard categorization scheme of the objectives in statistical learning: prediction/inference and regression/classification (Hastie, Tibshirani, and Friedman 2008; James et al. 2013). For the above reasons, the following argument is to a great extent tailored towards those readers who did not use the technique so far.

But the article is neither a manual on neural networks. Due to space limitations and envisaged goals, it is beyond the scope of this paper to discuss here in detail issues related to model specification and training. Readers interested in the relevant details are encouraged to consult the cited literature. Yet, some fundamentals in neural networks conceptualization and operationalization are tackled, since it is assumed that the social science community is still not too familiar with some basic neural concepts. Here, for the sake of clarity, only rudimentary remarks should suffice. This also gives an opportunity to introduce nomenclature.

Neural networks are computational models inspired by a biological formula. There are dozens of known types of NNs (= ‘architectures’), but for current purposes, the ‘classic’ network should be illustrative. Figure 1 presents the first formal representation of a biological neuron (McCulloch and Pitts 1943).

¹ Hereafter, neural networks (NNs). It is the shortened and well-recognized version of the traditional term. Admittedly, neural networks are just one of many machine learning methods. At the same time, however, NNs may well serve as a platform for discussing profound conceptual challenges.

Figure 1. The McCulloch-Pitts model of a neuron.



Source: (Rothman 2020).

The model building blocks are: inputs (= independent variables), weights, and output (= dependent variable). What the model does is processing the data from inputs I_i , $i = 1, 2, \dots, N$ to output y . The input data are aggregated with the respective weights ($I_i \cdot W_i$, $i = 1, 2, \dots, N$) and a non-linear functional transformation is applied to result in setting the y value.

The above formula is expanded to a complex structure that consists of layers of neurons. Such a multilayered structure is called a multi-layer perceptron (MLP) (Rosenblatt 1961) that is used as a workhorse in many analyses up to today. Actually, the connectivity between the layers makes NNs a viable option to reach for when dealing with many nonlinear problems.

It is also important to acknowledge that the model features do not aim at direct mirroring of biological nerve cells. It is rather an inspiration through applying their strengths: stable, resistant to damage, nonlinear, parallel, and generalizing signal processing. This translates into the applicability of NNs models to data processing without thorough investigation of the theoretical background of the investigated problem, its governing laws, or causal links (functional relations). In many real-world situations, we have no knowledge of these features, which severely hinders traditional modelling. The problems that occur in political science are no exception; quite the contrary – this is too often normal to be ignored. Thus, NNs may be considered in the social sciences toolbox as a workable solution in more cases than not.

But how does the network work? Without going too deep into technical jargon and mathematical notation, it may suffice to say that the objective of the network is to set the most optimal values of weights W_i , $i = 1, 2, \dots, N$. Here, the most optimal values mean the least error value as defined by the difference between the input pattern(s) and network's output calculated with dedicated algorithms. Each round of calculations is called an 'epoch.' The process is known

as ‘learning’ of the network, since the model *learns* the structure of the data. Let us turn to a real-life illustration for clarity. Children, when presented with new information, have two options to choose from: either they will learn it by rote or they will *understand* the problem. Neural networks are rather the latter, since they are designed to represent data structure parallelly when presenting new data, i.e. to understand the problem at hand. This single feature itself makes NNs extremely open to process incoming information that was not known before, i.e. the network is able to *learn* but not necessarily to *memorize* data structure (Schrodt 2000:305). Consequently, the knowledge gained is generalizable and can be extended to new tasks. The question of network learning is also tackled below when the categorization scheme is discussed.

Data is another critical issue. They are typically *somehow* correlated and not normally distributed, there are missing values, noise and/or outliers – to name a few features important enough to be ignored.² Hence, the prospect of building more robust models should be even more valuable, with neural networks being one of intriguing candidates for considering. This would serve as another argument for showing some interest in the NNs; the following review illustrates how the potential may be realized. Notwithstanding the above, data used for model training are usually divided into two separate clusters: training set and test set. The first is used for model specification, whereas the second is used for assessing its performance. Normally, as a rule of thumb, 75%-25% data allotment is practiced, respectively.

As was already mentioned, there are many NNs architectures with the McCulloch-Pitts model being a pioneer. Later developments have their usual ebbs and flows of interest, but the digital revolution made neural modelling one of the most successful cover stories (occasionally, literally) since decades. Deep learning deserves its credit here. It is successfully applied in data-rich and highly nonlinear problems such as speech and pattern recognition, anomaly detection, multi-label classification, computer-assisted translation, bioinformatics, etc. Here, the learning algorithm is used to find new patterns in data and harness knowledge to new (incoming) observations (Hinton and Salakhutdinov 2006). Until quite recently, due to the number of computations involved, deep learning approaches were tantamount to using the most powerful computers. Fortunately, there are some developments that help overcome such a limitation. One of the most promising areas is the so-called transfer learning that allows for the use of models already tuned to similar, yet new problems (Ruder 2017). Another advancement are regularization techniques such as dropout, which is briefly discussed in the closing section of the article. Since many of the above issues may seem to be rather advanced to the novice, the reader is strongly advised to consult the referenced literature.

Acknowledging the above details, the question of NNs applicability in political science must be asked. Are neural models more robust than other approa-

² Indeed, as data and datasets are critical factors in machine learning and neural networks, their review calls for a separate discussion.

ches? Do they provide any new insight to already practiced research? Are NNs able to address issues that used to be perceived as problematic? And finally, are they cure for all the methodological problems at hand?

It is claimed here that, on the most general level, a number of properties of neural networks bear reference to social phenomena. This article, then, is an attempt to explain why NNs may be relevant for some theoretical and empirical puzzles in political science.

However, before we proceed, one clarification is pending: a categorization scheme. As already mentioned, NNs are linked to signal/data processing. It is customarily applied to issues in the generic category of 'pattern recognition' and has two meanings. First, literally, it is about the analysis of patterns as graphic representations of certain phenomena (pictures, videos, etc.). Second, figuratively, it refers to patterns as structural elements that describe the phenomenon under investigation. In this paper, the focus is on the second meaning of the term, dealing with the classification of patterns in political behavior of individual and collective actors. However, commonly used, the term 'pattern recognition' itself is quite vague for at least two reasons. First, there are different 'patterns' addressed with NNs: some of them are qualitative and some may well be quantitative. Second, 'recognition' involves a gamut of possible analytical approaches. Both issues imply critical decisions on different architectures for different data and research problems. Thus, how to structure the following argument?

Like in any typology task, the key issue is related to the applied criteria. Some examples are included in Table 1.

Table 1. Examples of categorization approaches in neural networks applications.

criteria	categories	applications
learning method	supervised	regression/classification
	unsupervised	clustering
	reinforced	control
output variable(s)	quantitative	regression
	qualitative	classification
learning objectives	estimating functional relations	prediction vs. inference
	probability assignment	regression vs. classification

One of the most common classification designs is based on the widely accepted assumption that the learning method matters. Consequently, three main categories are assigned: supervised, unsupervised, and reinforced learning. For reasons related to the subject of our interest, further consideration will be limited to the first two types:

In supervised learning, the goal is to predict the value of an outcome measure based on a number of input measures; in unsupervised learning, there is

no outcome measure, and the goal is to describe the associations and patterns among a set of input measures (Hastie et al. 2008:xi).

In other words, a supervised model is expected to determine the level of similarity of input data with expected values (patterns). This may easily be illustrated in classification design: the network is provided with already determined classes and is tasked with declaring whether new observations belong to any of the known categories. If the answer is negative, such instances are treated by convention as outliers. Let us now return to the learning algorithms that were previously mentioned and used to set the model parameters. One of the most commonly used is the *backpropagation* algorithm that is considered a gold standard in MLP models (Werbos 1974, 1994).³ Here, first, random weights values are assigned, and then in the network learning process, errors from the last layer are ‘propagated back’ (hence the name) to earlier layers to modify the weights accordingly. The ultimate aim is to minimize the so-called mean square error: It is an aggregated sum of errors of the last neurons in a network. By changing the weights values, we can optimize the performance of the models in a series of epochs.

The other logic applies to unsupervised learning. Here, the network is not presented with any predetermined expected outcomes and therefore the model is set free to determine the outputs, there are no “measurements of the outcome” (Hastie et al. 2008:2). In our classification task example, this would mean that the input data are structured based on their statistical performance to set dimensions accordingly.

The second generic categorization criterion – characteristics of output variable(s) - however, bearing a different name, has quite similar logic: the model serves as either a ‘value generator’ or ‘category detector’. But here what matters most are variables’ features: whether they are, respectively, numerical or categorical (James et al. 2013:28–29).

Another proposal pertains to objectives of statistical learning: either functional relations or class probabilities are our research focus. In the first case, we pertain to two research designs: prediction or inference. The difference between them is intuitively approachable: prediction relates to problems where setting numerical values to dependent variables is critical, whereas inference assumes that our concern lies in obtaining knowledge about the way dependent variable is affected as independent variables change. So, prediction objective is to understand how dependent variable changes as a function of time, whereas in inference, as a function of features of independent variables.

The second learning objective, i.e., class probability assignment, pertains to regression vs. classification dichotomy. Here, one delves into the question of estimating class probabilities: if a given observation may be assigned a numerical value than it is a regression-type research design and when observations are allotted to categories, then we deal with a classification task.

³ Interestingly, the algorithm was originally applied in political science research.

2. NEURAL NETWORKS APPLICATIONS.

The above comments of the categorization scheme allows for some discussion since it seems that many relevant conundrums may stem from setting disciplinary boundaries rather than methodological boundaries. At least one question arises: How to find a middle ground in the above bewildering world between state of the art literature and our main goal here, i.e. applicability of NNs in political science? Since the article is basically addressed to political scientists (and practitioners in neighboring disciplines), their perspective is deemed to prevail. Therefore, the categorization approaches in Table 1, although rather ‘technical’, were out of scope here. And the literature review discussed further shows that notwithstanding technicalities, several application domains may be revealed. Out of a myriad of available categorization schemes, the one applied below was based on the nomenclature of the Research Committees of the International Political Science Association (IPSA).⁴ Obviously, out of 50 current committees only some of them are amenable to neural networks applications; these examples structure the following argument under respective subsections. Strictly speaking, the literature review was developed based on searches in the Scopus, Web of Science, and EBSCO databases. The search criteria were limited to peer-reviewed papers published in English in political science and international relations between the years 1991-2019. Books and chapters of books were also considered after reviewing the literature (citations). The analysis ruled out texts that treated political phenomena only as examples of applications of specific algorithmic solutions. Therefore, cases were included only if NNs were used to solve specific research puzzles in political science and international relations but not in STEM disciplines.

The criteria adopted above are obviously rather conservative and, as with every arbitrary decision in scientific debate, arguable. Many conference papers, post-conference publications, or (particularly doctoral) dissertations were omitted if they failed to meet the above criteria. Each narrowing of the search area is artificial, but this seems to be the necessary cost of a research design with a clearly defined objective and rational standing to support these objectives. The findings are discussed in detail below.

2.1. Concepts and Methods.

Due to the main research objectives of the article, the category based on Concepts and Methods IPSA research committee was one of the most troublesome, since for obvious reasons any research related to methods and techniques would fall into its focus. This would result in putting all applications into one category. However, it is acknowledged that only these examples that pertain to

⁴ Being the largest and the oldest professional association, IPSA’s research committees rather comprehensively cover research areas in political science as practiced across the globe.

specific methodological/conceptual puzzles are relevant here. Therefore, most prediction/forecasting research applies. Before discussing the same applications, let us make a few introductory remarks.

Neural networks, despite their known limitations, when compared to other techniques show relatively greater utility, especially in the case of non-stationary time series. This is because more classical approaches, such as the autoregressive-moving-average model (ARMA) and its variants, impose the need to ensure the stationarity of the time series. Such a feature considerably facilitates the analysis; if the regular structure of data is known, it is relatively easy to indicate future values. This assumption, however, is only exceptionally met in empirical data, hence ARMA models use various techniques (most often differencing) to bring the series to stationarity. NNs do not impose such a restriction, so they are used to study data in which seasonality or trend are not easy to detect and control.⁵

The usefulness of a neural approach is also evident in the case of multivariate time series. Here, the variable of interest depends not only on its historical values but also on the other related variables. Data that describe social phenomena often have such a complex structure. A clear illustration of this may be the process of predicting election results. It would be difficult to reasonably assume that the *only* variable that affects the future election result is the previous outcome. Indeed, the fact that a candidate/political party/coalition has previously won the election is not a guarantee of future success.

In light of the above comments, it should be clear that the use of NNs in time series analysis does not entail establishing functional relations describing the phenomenon under analysis. This is part of the methodological toolbox of traditional statistical analysis that prevails in econometric methods.

Let us turn to some examples of the use of NNs in time series analysis. Publications that belong to this group cover mostly elections, for example, in the Netherlands (Eisinga, Franses, and van Dijk 1998), Great Britain (Borisjuk et al. 2005) or the USA (Heredia, Prusa, and Khoshgoftaar 2017). For illustrative purposes, let us shed some light on the Borisjuk *et al.*, 2005 research. Their objective was to predict election final results based on the UK history between 1835 and 2001. The neural model consists of 11 input variables based on a survey of expert assessments of the political situation. Their closed responses (Yes-No) were coded as a binary variable (1-0, accordingly). The output was also binary: 0 for the success of the opposition party, 1 for the incumbent. The network was able to predict the real winner with 71-91% accuracy.

Some readers may be likely to think that a research design like this is highly speculative. After all: How is it possible to predict anything in the area where our expectations are so far from determinism, and the belief in the open and competitive nature of the electoral procedure is one of the pillars of theoretical,

⁵ Additionally, in the case of time series with seasonality, traditional approaches assume the cycles have a symmetrical structure, whereas in the case of NNs it is possible to model their asymmetry.

philosophical, and empirical features of democracy. Also, some methodological arguments pointing to poor data quality or their insufficient number have been raised in a critical tone (Linzer 2013; Ulfelder 2012; Wang et al. 2015). However, it turns out that this area may be the subject of systematic scientific research. All things considered, what we have at our disposal is a major political phenomenon characterized by regularity, relatively well-recognized theoretical framework, and availability of data from a variety of sources, e.g. opinion polls and electoral register data. Importantly, the analysis of data from 85 countries indicates that, with the growing body of available data, it is possible to be more accurate in predicting election results (Kennedy, Wojcik, and Lazer 2017).⁶

The basic argument related to the limitations of time series analysis is the availability of *appropriate* data. In this context, there are concerns that we still deal mostly with observations registered only in the annual interval and at the level of countries rather than smaller territorial units (Spiegeleire 2016:9).

Apart from the electoral procedure, some methodological puzzles in prediction/forecasting are also discussed in the article on anticipation of the voting results in the US Congress (Khashman and Khashman 2016).

2.2. Security, Conflict and Democratization/Quantitative International Politics.

Research in these two areas is burgeoning and the most studied subcategory is the study of armed conflicts (Brandt, Freeman, and Schrodt 2011, 2014; Cederman and Weidmann 2017; Cranmer and Desmarais 2017; Schneider, Gleditsch, and Carey 2011; Schrodt 2004). Because of many detailed methodological ramifications, a closer look would be beneficial.

The first analyses conducted in this field were based on traditional statistical techniques such as linear regression models, yet these failed to address the statistical characteristics of variables; for example, lack of a normal distribution or the problem of nonlinearity (Beck, King, and Zeng 2000:22, 24; Diehl 2006:207). A separate problem is the questionable effectiveness of evaluating traditional prediction models using metrics such as the root mean squared error (RMSE) or the mean absolute error (MAE). The main objection is the sensitivity of these measures to outliers (Brandt et al. 2011, 2014).

The commonly used data structures are ‘dyads’, i.e. sets of two elements, in this case pairs of countries, that are in interaction. However, the question remains: Which countries (dyads) should be included? Should we incorporate all pos-

⁶ See the same source for a discussion on the limited importance of opinion polls and structural factors (e.g. economic indexes) in the accurate prediction of election results. Cf. (Esfandiari et al. 2012) More examples of the usefulness of forecasting with the use of other methods in the presidential and Congressional elections in the US see e.g. (Campbell 2017). In the USA, history of using forecasts in the media outlets goes back to the presidential elections of 1952 when the UNIVAC I computer was used to work out the real time forecast for the purposes of CBS Evening News (Morris 2019).

sible countries or only those whose choice is justified by the adopted theoretical assumptions? In addition, the data is extremely diverse in terms of frequency. It turns out that conflict situations belong to a much smaller group of cases than peaceful situations (Beck et al. 2000:22, 23–24, 28). This results in a number of significant methodological reservations addressing the question of data heterogeneity (Chio and Freeman 2018:57–58; Japkowicz and Stephen 2002; King and Zeng 2001a, 2001b; Muchlinski et al. 2016; Schneider et al. 2011).

The studies discussed present a considerable intellectual challenge not only because of the complex relationships between independent variables and questionable data quality. Another problem which has been ignored for many years is the static nature of the proposed models. It has been customary, for example, to ignore the fact of transformations occurring in international relations, of which territorial changes are the most obvious (Cederman and Weidmann 2017:475). Only more recent work has attempted to remedy these limitations. In order to illustrate the issue, let us examine an example of a relatively simple research strategy (Schrodtt 1991).

Schrodtt's analysis was fueled by a database that contained 310 cases of international armed conflict that occurred between 1945 and 1974. These were described with 47 variables. The input layer of the model included 24 independent variables and the output layer three dependent variables that describe the probability of conflict. The comparison of the neural network results with other, more traditional techniques demonstrated that NNs turned out to be only subtly more effective, which Schrodtt links not so much to the limitations of the networks but to the data quality.

Some of these difficulties came to light at the beginning of the millennium in an interesting exchange of views published in the *American Political Science Review*. Notwithstanding its 20-year history, let us take a closer look at the presented arguments, since they well illustrate many of the NNs intricacies.

The title of the first article expressed the reservation that the authors were only concerned with: a *conjecture* (Beck et al. 2000).⁷ Despite such a 'prudential' approach, the arguments were strong enough to have become a subject of heated debate in the scholarly community. The starting point of their analysis was dissatisfaction with the state of research on armed conflicts, referring to the models used at the time by the term "fragility" (Beck et al. 2000:21). Their main complaint was about the poor prediction performance of the models. Furthermore, the forecasts were based on the database that was used for the specification of the model (in the sample), while its evaluation should have been carried out using new data (out-of-sample) (Beck et al. 2000:21–22).⁸ This strategy, known

⁷ Hereafter BKZ 2000.

⁸ To put it in other words, it is one thing to tune the model to the known data and another to the future data. Some argue that formally statistically significant variables do not necessarily improve the model quality, since this may be related to overfitting the training (in-sample) data set, hence the proposal of the use of machine learning and a test set (Ward, Greenhill, and Bakke 2010). The argument is widely accepted in literature (Colaesi and Mahmood 2017; Cranmer and Desmarais 2017; Hindman 2015; Muchlinski et al. 2016; Ward and Beger 2017).

as cross-validation, guarantees the ability to generalize acquired knowledge and transfer it to new cases while also protecting against overfitting (Berk 2006; Colaresi and Mahmood 2017; Ward and Beger 2017; Witmer et al. 2017). BKZ 2000 research was based on a collection of over 23,500 cases of dyads-years, for the period from 1947 to 1989. The dependent variable was dichotomous: It takes the value '1' in the case of a militarized interstate dispute and '0' in the opposite. The independent variables include: neighbors, participation in alliances, similarity in the assumptions of external policy, asymmetry of military strength, degree of democracy, and the number of years since the last conflict between the countries in the pair. This data set was then divided into an in-sample training set with observations for 1947-1985 and a test set for 1986-1989. Finally, the analysis involved a network with 25 hidden nodes and a logistic activation function. The resulting model correctly predicts almost 25% of conflict situations and almost 100% of situations without conflict (the traditional logistic model predicts 0% and 100% respectively, i.e. it is useless). In effect, the authors concluded that the solution they had proposed was a "superior statistical model of international conflict" (Beck et al. 2000:32). This set the stage for some criticism, which was expounded in the article published four years later.

A critical response was provided by the authors who provocatively entitled their text 'Untangling Neural Nets' (de Marchi, Gelpi, and Grynaviski 2004).⁹ To begin with, the authors questioned the data used by BKZ 2000 as tendentially selected from the body of research, in order to demonstrate the superiority of neural networks over the logistic regression model. Additionally, the value of the results questioned was challenged due to the fact that NNs introduce "unnecessary uncertainty about causal relationships because of inefficient estimates" and "precludes hypothesis testing that focuses on measuring or comparing the effects of particular variables" (de Marchi et al. 2004:372). In this way, MGG seemed to favor the traditional logic of scientific research: theoretical assumptions → hypotheses as to functional relations → verification of hypotheses.

In fact, MGG modified the BKZ 2000 proposal by including three re-examined variables and adding a new variable: the country's status as a superpower (according to the theoretical assumption that superpowers show a greater propensity to use force). The added methodological value was the postulate to use the ROC (receiver-operating characteristic) curve, which interpretation is relatively simple and does not require any arbitrary probability level to accept or reject the model quality.

Based on these assumptions, MGG proposed four models: two neural networks, a logistic model, and a linear model based on discriminant analysis. The analysis demonstrated that all four specifications were of similar predictive performance. Moreover, for the test data (1985-1989) linear and logistic models showed even better properties; the authors explained this in terms of the inclusion of the modified variables. As the authors briefly concluded, "We are not aware of any theories in the study of international conflict that specify a func-

⁹ Hereafter MGG.

nal form so complex as to require, or even suggest, a neural network” (de Marchi et al. 2004:378).

The response provided by Beck, King and Zeng¹⁰ confirmed that the model quality had to be verified based on the ROC curve and out-of-sample data set and not the data used in model specification (Beck, King, and Zeng 2004:379). Furthermore, they pointed out that the use of both logistic models and neural networks resulted in coefficients not being interpreted directly. This is an argument that gains importance if we consider that NNs have more parameters than logistic models.¹¹ Also, BKZ 2004 maintained the argument on the flexibility of neural networks in the analysis of the phenomena whose functional form is not conclusively theory-defined. Specifically, the authors respecified one NN and the two logistic models. Based on the ROC curve, BKZ 2004 demonstrated that the neural network possessed a much greater ability to predict the dependent variable. Much information is provided, in particular, by the analysis of the probability surface graphs of the outbreak of war in relation to the degree of democracy in a given country. However, the question remains: Is this feature resulting from the fact that the two countries are democracies or is it because they are similar? The analysis of the probability surface graph shows that it is the similarity of political systems that seems to account for more peaceful co-existence rather than the two neighboring countries being democracies. This is an important and epistemically interesting conclusion, which was only possible with the use of neural networks.

But being too optimistic is a little premature. The main dilemma resulting from the debate can be expressed as follows: research taking advantage of neural networks, like any other research, must be preceded by a specific theoretical reflection which provides a framework for scientific investigations. So far, this statement has been relatively far from controversy for most, if not all, researchers, including those working within the social sciences. However, in this context, a significant conceptual problem arises. If an assumption of a complex (nonlinear) structure of the investigated phenomenon is adopted, is it acceptable to refer to linear techniques? Beck, King, and Zeng answer is negative, while that of de Marchi, Gelpi and Grynawski is positive.

Another example in this category is the work that analyzes international conflicts from the mid-19th century to the early 1990s (Lagazio and Russett 2003).¹² The model specification clustered independent variables into two categories according to the two classic theoretical approaches in international relations: realism and liberalism. These are alliances, neighborhood, geographical distance, being a superpower, balance of power (realism) and level of democracy, economic interdependence and membership in international organizations (liberalism). The authors focused on the extent to which these variables influenced the con-

¹⁰ Hereafter BKZ 2004.

¹¹ On the other hand, there are techniques that reduce the number of parameters in NNs (Beck et al. 2004:381).

¹² Cf. (Lagazio and Marwala 2006; Marwala and Lagazio 2011)

flict/peace situation (dependent variable) and concluded that along with stronger economic interdependence, the tendency to be involved in conflict decreased. However, aggressive moves were found to be accompanied by a decline in the level of democracy.

Better recognition of the probability of external or internal conflict in a given place in a given year has been mostly due to the development of big data studies. In fact, computer analyses of large data sets have opened entirely new research horizons. One of such developments aims at adding new variables to the models, for example, including mass media coverage of conflicts (Chadefaux 2014; Leetaru 2013). These kinds of search for behavioral patterns can be found under the deliberately coined term, i.e. *culturomics* (Leetaru 2011). One of the early examples includes the analysis of data from the database of digitized books, published in 1800-2000, using the frequency of words and expressions (Michel et al. 2011). Let us also mention the GDELT news monitoring project with its historical archives that date back to 1979. The ‘density’ of its information resources (a huge amount of data reported in real time) provides greater analytical possibilities compared to structural variables, such as the level of democracy in a country (Cederman and Weidmann 2017:475). Another innovation is an attempt to model the location of conflicts down to the level of country-specific administrative units (Weidmann and Ward 2010).

Finally, let us end the part on studying armed conflicts with recent attempts to use machine learning techniques. Their applicability has gained recognition since a foundational study was published three decades ago (Schrodt 1990). Therefore, their use in political science calls for a separate study, yet here only some introductory remarks are provided. Strictly speaking, although this paragraph is based on a study that has not involved NNs, the two approaches are closely linked. To be precise, the issue of conflict prediction can be treated as a fairly typical example of a problem within the framework of supervised learning. Therefore, according to the classical enunciation of machine learning, the objective of the analysis (*T*) is to develop forecasts for future cases based on the available data (*E*), using the applicable accuracy measures (*P*).¹³ The main advantages of this are twofold. First, modeling phenomena within the framework of supervised learning requires the breaking down of data into three different sets (learning, test, and validation) to avoid overfitting. This postulate, as we know it, is neither new nor particularly trivial. It even gains in importance when combined with the second advantage, which is the possibility of avoiding (in the same research procedure) the danger of specification of the model that is not generalizable. In other words, we are caught between two extreme situations of overfitting and underperformance, while machine learning provides a promising attempt at finding balance (Colaresi and Mahmood 2017:198–99). If the logic of the research process (i.e. iterative analysis of the subsequent versions of the model) is added to the above, it turns out that the benefits of machine learning might be worth considering in future studies of predicting armed conflicts. Also, other successful

¹³ The notation corresponds to the classical definition of machine learning (Mitchell 1997:2).

political science applications based on machine learning are worth acknowledging (Cranmer 2019).

Another example of a research area in which neural networks have been applied is the analysis of factors that affect the collapse of states (King and Zeng 2001a). Here, the dependent variable is dichotomous: the fall of a state in a given year is coded as “1”, and its survival as “0”. On the basis of theory, an assumption was made about the mutual links and non-linear nature of the relationships between variables. Six input variables were included for the years 1955-1998 and 195 countries: the degree of democracy, commercial openness, infant mortality, proportion of population in the armed forces, population density, and effectiveness of legislation. Quality assessments were made again using the ROC curve and neural networks have proven to be more effective than traditional techniques.

Other examples of research in this category include: democracy and democratization patterns (Buscema, Ferilli, and Sacco 2018; Buscema, Sacco, and Ferilli 2016; Kimber 1991; Vanhanen and Kimber 1994), military expenditure (Refenes, Kollias, and Zapranis 1995), political analysis of the peace processes (Ballén 2014), global evolution of political systems (Buscema, Ferilli, and Sacco 2017), systemic risk and insecurity (Cimpoeru 2015; Grigora 2013), peace mediation (Khashman and Khashman 2017), political violence and instability (Blair, Blattman, and Hartman 2017; Goldstone et al. 2010; Iswaran and Percy 2010; Rost, Schneider, and Kleibl 2009), and armaments and rivalry between states (Andreou and Zombanakis 2001, 2011).

2.3. International Political Economy (IPE).

IPE research is represented here by two major subfields: sanctions and globalization issues. A text on the effectiveness of sanctions may be illustrative here (Bearce 2000). The independent variables (network inputs) consist of 11 neurons that characterize the phenomena defined by the theoretical assumptions found in the literature. Among the examples of such variables are the year in which the sanctions were introduced (hypothesis: a decreasing effectiveness of sanctions with the progress of globalization) or international help for the country to which sanctions were applied (hypothesis: the existence of international aid programs reduces the effectiveness of sanctions). The model also takes into account the threshold node; an ‘excessive’ element without empirical equivalent, which corresponds to a constant term in classical statistics. The output of the network is only one neuron, which marks the measure of the effectiveness of sanctions expressed on the ordinal scale. The data set includes 115 cases of sanctions gathered by the Institute for International Economics. The network results have been compared with the ordinary least squares (OLS) method and the probit model. In both cases, neural networks turned out to be more accurate.

The article on globalization (Dorado-Moreno, Sianes, and Hervás-Martínez 2016) supplements the IPE category.

2.4. Administrative Culture/Public Policy and Administration.

Examples of research on administrative arrangements include three contributions that analyze various facets of the quality of state mechanisms. The first paper studies factors related to e-government services in Egypt (Mostafa and El-Masry 2013). The authors applied three NN architectures and three other machine learning approaches to investigate the influence of independent variables. The second article focuses on measuring the risk of corruption of civil servants using political party affiliation data (Carvalho et al. 2014). Interestingly enough, methodologically it was also based on the machine learning approach. Specifically, the authors used four machine learning classification tools: Bayesian networks, support vector machines (SVMs), random forest, and backpropagation NNs. Data mining metrics allowed for comparison of a specified model with human expert assertions. No major discrepancies were found. And the third item in this category also investigates the problem of public corruption (López-Iturriaga and Sanz 2018). Here, self-organizing maps were used to build an early warning anticorruption detector. The results show that economic variables mostly contribute to the vulnerability of corruption.

As the above paragraph indicates, from the statistical learning point of view, all three contributions fall into the classification category. But what makes them even more valuable is the fact that the authors juxtaposed standard NN architectures with other machine learning techniques. Recognizing their strengths and weaknesses, it turns out that neural networks are worth considering as a viable option when classification is at stake.

2.5. Political Communication.

Formally, here classification designs are also most often applied. Specifically, this refers to the classification of content in political discourse. Applications include radial basis NN (Ran 2015) and multi-scale convolutional neural networks (Bilbao-Jayo and Almeida 2018). The first article builds a ‘discourse classifier’ in Spanish, Finnish, Danish, English, German, French, and Italian. The authors improved the model by introducing two additional features: the previous sentence in the text corpus (party manifestos) and the political leaning of the speaker. This results in a significant improvement in the classifier.

Another approach was introduced in the article on the detection of political ideology detection (Iyyer et al. 2014). The authors used a recursive neural network to identify ideological bias in a sentence-level political text. Specifically, gold standard ‘bag of words’ classifiers that ignore the linguistic context of the narrative were substituted with deep learning approaches applied in sentiment analysis. This allows for discovering compositional effects in a text leveraging analytical applicability to ideological bias detection.

A different architecture (convolutional neural network) was implemented in a paper on electoral violence and unstructured social media text (Muchlinski et

al. 2020). The rationale for utilizing this particular type of NN was its performance in computational linguistics and it was compared to the SVM baseline model. The three accuracy metrics used – precision, recall, and their harmonic mean, F1 – showed the superior classification performance of the neural network that is more accurately estimating electoral violence on social media.

There is at least one important extension to the above approach: a two-tier classifier that “uses social media data to identify collective action events occurring in the real world” (Zhang and Pan 2019:4). Here, although the convolutional neural network was also used, it was implemented to image classification whereas a combination of convolutional and recurrent neural networks with long short-term memory was used for textual analysis. Based on a data set that was built on nearly 140,000 collective action events, the authors were able to suggest that the neural-based classifier was at least as accurate as other classification systems.

2.6. Comparative Public Opinion.

This research starts with a standard model, i.e. feedforward architecture with backpropagation algorithm (Monterola et al. 2002). This simple but useful tool was used to classify undecided respondents to investigate their opinions and sentiments. The authors were quite cautious in interpreting their results, and one specific point is worth mentioning here: neural networks are “notoriously highly problem-specific” (Monterola et al. 2002:227). To put it in informal but evocative language, NNs are not one-size-fits-all solution: When trained on a particular dataset, they are tuned to the data structure and research question at hand.

Another example investigates survey data to predict voter turnout (Weber et al. 2018). Here, the neural network was used to model data from the European Social Survey for Germany. The trained network was evaluated against linear regression and two popular machine learning algorithms: SVMs and random forests. It turns out that NNs may be “capable of analyzing human decision patterns” (Weber et al. 2018:586). Interestingly, when referring to obstacles, the authors point out not the limitations of NNs but rather the question of data quality and accuracy.

The work on the ideological and political status of undergraduates complements public opinion research (Zhao, Li, and Li 2015). The authors reached for the backpropagation MPL and fuzzy evaluation method. Unfortunately, the article does not deliver details on the surveys used; it may be implied that results are valid for only one country thus formally they are not applicable to comparative approach.

2.7. Comparative Public Policy.

The last of our categories derived from the IPSA Research Committees nomenclature belongs to the one sparsely represented by NNs applications. Consistently with the already discussed criteria, only one contribution seems to be illustrative here (Hegelich 2017). Here at least two points are worth considering. First, the research design includes one of the machine learning approaches: deep learning. This implies a viable option for representing complex – i.e. nonlinear – prediction space, and thus may boost model performance. Second, the approach is used to investigate one of the seminal theoretical arguments in public policy research: Punctuated Equilibrium Theory (PET). Specifically, the author discusses similarities between neural networks and the core theoretical arguments of PET. Furthermore, the empirical part demonstrates that the deep learning model can contribute to academic research on one of the key PET assumptions: the relation between attention allocation in the policy system and major budget fluctuations (“punctuations”).

3. CONCLUSION

Recent developments in machine learning led to a significant reevaluation in the understanding and implementation of research design. This is true for just one approach covered here – neural networks. And against the background of other social sciences, political science is no exception. There are many ways to approach the issue and focusing on just one of the above areas would be a feasible option.¹⁴ However, here a more ambitious goal was perceived: to scan political science research across subfields in terms of possibilities—and perils—of using certain analytical technique. The analysis lets for some conclusions.

First of all, the reader should not be left with an impression that neural networks – like any research method – are a ‘one fits all’ approach. Thus, a few words of caution are in order. First, many NNs models are ‘data hungry’. Applying them to a research agenda with a few hundred observations is rarely workable. Consequently, any attempt to increase the size of the data set is a must; one of the most obvious suggestions is to broaden the time horizon or the number of variables. In other words, it is rather common to apply NNs to problems that are operationalized with thousands or more data points. This makes the family of neural networks and their relatives a good choice for research designs aimed at using ‘streaming data’, i.e., data that are continuously coming to tune the model. One of such applications is ‘outlier detection’ research that has some interesting history, mostly outside of political science. Consequently, it is intriguing to look for further applications, especially in the light of the fact that

¹⁴ Customary acknowledgements are due to one of the reviewers for pointing out this issue.

problems with outlying (abnormal) observations are not a novel issue in political science.

This leads to another, but related issue: to a much extent neural networks models are ‘garbage in-garbage out’ architectures. This means that the data used for modelling must be preprocessed. As we already know, data may be concomitantly interval and categorical (quantitative and qualitative) in a given architecture. Data may also be defined by obscure patterns. However, this does not mean that the data may be raw, i.e. unprocessed. Normally, input vectors must be of a specified form to deliver what is expected; this usually means standardization of variables (Garson 1998:91–94). Without preliminary study of data at hand, one is confronted with a possibility of obtaining a nonsense. Again, this sounds like a refrain to any research method, making a researcher using neural networks less a warlock and more a craftsman.

Criticism acclaimed, it seems that the discussed examples confirm the hypothesis presented at the outset of these considerations: neural networks may be useful in solving research problems in political science if only the following criteria are satisfied: (1) there is a scarcity of conclusively specified models over which fundamental theoretical debates have been settled, (2) there are complex and/or unknown functional features without well-defined data-generating mechanisms, (3) issues are operationalized with data having characteristics that hinder their traditional analysis, and (4) there are obstacles in collecting appropriate data in terms of their quantity and quality. As was shown above, political science is not particularly absent on issues defined by such characteristics. Neural networks are therefore one of the tools to reach for to deal with the extremely complicated social world when not only explanation but also prediction is at stake.

This argument is of special importance in light of another relentless conundrum: the black-box problem. To put it succinctly, for decades, it was impossible to interpret how neural networks arrive at the conclusions hindering inference. Indeed, nonlinear models and their results come with a price: parameters direct inexplicability (Faraway and Chatfield 1998:242–45; Zeng 1999:504, 2000:242–243). When they were mathematically manageable, their empirical understanding was not. Consequently, for many years, the only remedy was to use NNs with some reserve and caution in areas where knowing a generating mechanism was a primer; expert systems are one of such areas (Hayashi 1994; Knight 2017; Medsker 1994). Even the very nomenclature was contributing to the lack of transparency in neural networks: The commonly used term ‘hidden layer(s)’ is a case in point. Fortunately, recent developments overcome traditional concerns (Escalante et al. 2018; Ghorbani, Abid, and Zou 2017). One of such advancements are regularization techniques such as dropout which allows for pruning the inflated NNs architecture (Srivastava et al. 2014). This research field under the larger umbrella of ‘explainable Artificial Intelligence’ (Adadi and Berrada 2018) contributes, apart from other advances, to challenge the “reproducibility crisis” in science (Barber 2019; Helbing et al. 2017; Hutson 2018). Indeed, many research designs were ephemeral, that is, they were not able to be replicated, which is a gold standard in science (Prinz, Schlange, and Asadullah

2011; Raff 2019). Interestingly, the neural model closely follows its biological inspiration: The brain itself works according to the ‘black box’ metaphor, since many of its functions are still inexplicable (Castelvecchi 2016; Dreyfus and Dreyfus 1988:35). Consequently, there is a good reason to acclaim not only directly interpretable algorithms (Holm 2019). All in all, for many of practitioners, it is critical to get workable models without necessarily detailed knowledge behind them – just like for many skilled drivers it is not mandatory to know technical intricacies of car engine. On the other side of the spectrum are those who consider such an approach nothing but a “heresy.”¹⁵ Thus, from a practical point of view, one must acknowledge the limitations of neural networks, and one does not have to be overly phobic about them either. This would make NNs even less ‘hidden’ or inscrutable since we are stuck with similar quandaries when using *any* research tool/method. The three reservations mentioned above—data demand, data quality dependency and the black box conundrum—only shows that neural models are not universal cure for every problem and should be proceed with caution.

Considering the above arguments, it may be acknowledged that neural networks – alongside other machine learning techniques – have already secured their place in the political science toolbox, and consequently, the boundaries of the discipline itself will continue expanding. One of the possible extensions is the use of some of the existing approaches in a much more creative way. This would not mean a secondary analysis of the collected data, but a formulation of entirely new problems and the setting of new horizons. In fact, some pursuits seem to be already very promising (Colaresi and Mahmood 2017; Cranmer 2019; Hindman 2015). All this indicates that it is up to researchers whether the ‘brave new world’ will continue to be a curse or become a prevalent description of the state of the art political science in the future.

4. REFERENCES

- Adadi, Amina, and Mohammed Berrada. 2018. “Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI).” *IEEE Access* 6:52138–60. doi: 10.1109/ACCESS.2018.2870052.
- Andreou, Andreas S., and George A. Zombanakis. 2001. “A Neural Network Measurement of Relative Military Security: The Case of Greece and Cyprus.” *Defense and Peace Economics* 12(4):303–24.
- Andreou, Andreas S., and George A. Zombanakis. 2011. “Financial Versus Human Resources in the Greek–Turkish Arms Race 10 Years on: A Forecasting Investigation Using Artificial Neural Networks.” *Defence and Peace Economics* 22(4):459–69. doi: 10.1080/10242694.2010.539858.
- Ballén, Alfonso Toledo. 2014. “Political Analysis of the Peace Process in Colombia: A Simulation from Artificial Neural Networks.” *Visión Electrónica* 8(1).

¹⁵ Andrew Hunter quoted in (Garson 1998:16).

- Barber, Gregory. 2019. "Artificial Intelligence Confronts a 'Reproducibility' Crisis." *Wired*. Retrieved September 23, 2019 (<https://bit.ly/2mb6Es8>).
- Bearce, David H. 2000. "Economic Sanctions and Neural Networks: Forecasting Effectiveness and Reconsidering Cooperation." Pp. 269–95 in *Political complexity: non-linear models of politics*, edited by D. Richards. Ann Arbor: University of Michigan Press.
- Beck, Nathaniel, Gary King, and Langche Zeng. 2000. "Improving Quantitative Studies of International Conflict: A Conjecture." *American Political Science Review* 94(1):21–35.
- Beck, Nathaniel, Gary King, and Langche Zeng. 2004. "Theory and Evidence in International Conflict: A Response to de Marchi, Gelpi, and Grynaviski." *American Political Science Review* 98(2):379–89.
- Berk, Richard A. 2006. "An Introduction to Ensemble Methods for Data Analysis." *Sociological Methods & Research* 34(3):263–95. doi: 10.1177/0049124105283119.
- Bilbao-Jayo, Aritz, and Aitor Almeida. 2018. "Automatic Political Discourse Analysis with Multi-Scale Convolutional Neural Networks and Contextual Data." *International Journal of Distributed Sensor Networks* 14(11):1–11. doi: 10.1177/1550147718811827.
- Blair, Robert A., Christopher Blattman, and Alexandra Hartman. 2017. "Predicting Local Violence: Evidence from a Panel Survey in Liberia." *Journal of Peace Research* 54(2):298–312. doi: 10.1177/0022343316684009.
- Borisyuk, Roman, Galina Borisyuk, Colin Rallings, and Michael Thrasher. 2005. "Forecasting the 2005 General Election : A Neural Network Approach." *British Journal of Politics & International Relations* 7(2):199–209.
- Brandt, Patrick T., John R. Freeman, and Philip A. Schrodt. 2011. "Real Time, Time Series Forecasting of Inter- and Intra-State Political Conflict." *Conflict Management and Peace Science* 28(1):41–64. doi: 10.1177/0738894210388125.
- Brandt, Patrick T., John R. Freeman, and Philip A. Schrodt. 2014. "Evaluating Forecasts of Political Conflict Dynamics." *International Journal of Forecasting* 30(4):944–62. doi: 10.1016/j.ijforecast.2014.03.014.
- Buscema, Massimo, Guido Ferilli, and Pier Luigi Sacco. 2017. "What Kind of 'World Order'? An Artificial Neural Networks Approach to Intensive Data Mining." *Technological Forecasting and Social Change* 117:46–56. doi: 10.1016/j.techfore.2017.01.010.
- Buscema, Massimo, Guido Ferilli, and Pier Luigi Sacco. 2018. "The Meta-Geography of the Open Society: An Auto-CM ANN Approach." *Expert Systems with Applications* 99:12–24. doi: 10.1016/j.eswa.2018.01.017.
- Buscema, Massimo, Pier Luigi Sacco, and Guido Ferilli. 2016. "Multidimensional Similarities at a Global Scale: An Approach to Mapping Open Society Orientations." *Social Indicators Research* 128(3):1239–58. doi: 10.1007/s11205-015-1077-4.
- Campbell, James E. 2017. "A Recap of the 2016 Election Forecasts." *PS - Political Science and Politics* 50(2):331–38. doi: 10.1017/S1049096516002766.
- Carvalho, Ricardo Silva, Rommel Novaes Carvalho, Marcelo Ladeira, Fernando Mendes Monteiro, and Gilson Liborio de Oliveira Mendes. 2014. "Using Political Party Affiliation Data to Measure Civil Servants' Risk of Corruption." Pp. 166–71 in *2014 Brazilian Conference on Intelligent Systems*. IEEE.
- Castelvecchi, Davide. 2016. "Can We Open the Black Box of AI?" *Nature* 538(7623):20–23. doi: 10.1038/538020a.

- Cederman, Lars-Erik, and Nils B. Weidmann. 2017. "Predicting Armed Conflict: Time to Adjust Our Expectations?" *Science* 355(6324):474–76. doi: 10.1126/science.aal4483.
- Chadefaux, Thomas. 2014. "Early Warning Signals for War in the News." *Journal of Peace Research* 51(1):5–18.
- Chio, Clarence, and David Freeman. 2018. *Machine Learning and Security: Protecting Systems with Data and Algorithms*. Sebastopol, CA: O'Reilly Media.
- Cimpoeru, Smaranda. 2015. "Using Self-Organizing Maps for Assessing Systemic Risk. Evidences from the Global Economic Crisis." *Economic Computation and Economic Cybernetics Studies and Research* 49(2):71–89.
- Colaresi, Michael, and Zuhaib Mahmood. 2017. "Do the Robot: Lessons from Machine Learning to Improve Conflict Forecasting." *Journal of Peace Research* 54(2):193–214. doi: 10.1177/0022343316682065.
- Cranmer, Skyler J. 2019. "Introduction to the Virtual Issue: Machine Learning in Political Science." *Political Analysis* 1–9.
- Cranmer, Skyler J., and Bruce A. Desmarais. 2017. "What Can We Learn from Predictive Modeling?" *Political Analysis* 25(2):145–66. doi: 10.1017/pan.2017.3.
- Diehl, Paul F. 2006. "Just a Phase? Integrating Conflict Dynamics Over Time." *Conflict Management and Peace Science* 23(3):199–210. doi: 10.1080/07388940600837490.
- Dorado-Moreno, Manuel, Antonio Sianes, and César Hervás-Martínez. 2016. "From Outside to Hyper-Globalisation: An Artificial Neural Network Ordinal Classifier Applied to Measure the Extent of Globalisation." *Quality & Quantity* 50(2):549–76. doi: 10.1007/s11135-015-0163-7.
- Dreyfus, Herbert, and Stuart Dreyfus. 1988. "Making a Mind Versus Modeling the Brain." in *The Artificial Intelligence Debates: False Starts, Real Foundations*, edited by S. Graubard. Cambridge: MIT Press.
- Eisinga, Rob, Philip Hans Franses, and Dick van Dijk. 1998. "Timing of Vote Decision in First and Second Order Dutch Elections 1978-1995: Evidence from Artificial Neural Networks." *Political Analysis* 7:117–42.
- Escalante, Hugo Jair, Sergio Escalera, Isabelle Guyon, Xavier Baró, Ya mur Güçlütürk, Umut Güçlü, and Marcel van Gerven, eds. 2018. *Explainable and Interpretable Models in Computer Vision and Machine Learning*. New York, NY: Springer Science+Business Media.
- Esfandiari, A., H. Khaloozadeh, M. Esfandiari, and Z. J. Fard. 2012. "Machine Learning in Parliament Elections." *Research Journal of Applied Sciences, Engineering and Technology* 4(19):3732–39.
- Faraway, Julian, and Chris Chatfield. 1998. "Time Series Forecasting with Neural Networks: A Comparative Study Using the Airline Data." *The Journal of the Royal Statistical Society, Series C (Applied Statistics)* 47(2):231–50.
- Garson, G. David. 1998. *Neural Networks: An Introductory Guide for Social Scientists*. London: Sage.
- Ghorbani, Amirata, Abubakar Abid, and James Zou. 2017. "Interpretation of Neural Networks Is Fragile." Retrieved March 4, 2019 (<https://arxiv.org/abs/1710.10547>).
- Goldstone, Jack A., Robert Bates, David Epstein, Ted Robert Gurr, Michael B. Lustik, Monty G. Marshall, Jay Ulfelder, and Mark Woodward. 2010. "A Global Model for Forecasting Political Instability." *American Journal of Political Science* 54(1):190–208. doi: 10.1111/j.1540-5907.2009.00426.x.
- Grigoraş, Răzvan. 2013. "A Systemic Approach of Predicting the State of Security: Neural Modelling." *Revista Academiei Fortelor Terestre* 18(3):231–39.

- Hastie, Trevor, Robert Tibshirani, and Jerome Friedman. 2008. *The Elements of Statistical Learning. Data Mining, Inference, and Prediction*. Second Ed. Springer.
- Hayashi, Yoichi. 1994. "Neural Expert System Using Fuzzy Teaching Input and Its Application to Medical Diagnosis." *Information Sciences - Applications* 1(1):47–58.
- Hegelich, Simon. 2017. "Deep Learning and Punctuated Equilibrium Theory." *Cognitive Systems Research* 45:59–69. doi: <https://doi.org/10.1016/j.cogsys.2017.02.006>.
- Helbing, Dirk, Bruno S. Frey, Gerd Gigerenzer, Ernst Hafen, Michael Hagner, Yvonne Hofstetter, Jeroen van den Hoven, Roberto V. Zicari, and Andrej Zwitter. 2017. "Will Democracy Survive Big Data and Artificial Intelligence?" *Scientific American*. Retrieved September 13, 2019 (<https://bit.ly/2KxppiR>).
- Heredia, Brian, Joseph Prusa, and Taghi Khoshgoftaar. 2017. "Exploring the Effectiveness of Twitter at Polling the United States 2016 Presidential Election." Pp. 283–90 in *IEEE 3rd International Conference on Collaboration and Internet Computing (CIC)*.
- Hindman, Matthew. 2015. "Building Better Models: Prediction, Replication, and Machine Learning in the Social Sciences." *The ANNALS of the American Academy of Political and Social Science* 659(1):48–62. doi: [10.1177/0002716215570279](https://doi.org/10.1177/0002716215570279).
- Hinton, Geoffrey, and Ruslan Salakhutdinov. 2006. "Reducing the Dimensionality of Data with Neural Networks." *Science* 313:504–7.
- Holm, Elizabeth A. 2019. "In Defense of the Black Box." *Science* 364(6435):26–27. doi: [10.1126/science.aax0162](https://doi.org/10.1126/science.aax0162).
- Hutson, Matthew. 2018. "Artificial Intelligence Faces Reproducibility Crisis." *Science* 359(6377):725–26. doi: [10.1126/science.359.6377.725](https://doi.org/10.1126/science.359.6377.725).
- Iswaran, N., and D. F. Percy. 2010. "Conflict Analysis Using Bayesian Neural Networks and Generalized Linear Models." *The Journal of the Operational Research Society* 61(2):332–41. doi: [10.1057/jors.2008.183](https://doi.org/10.1057/jors.2008.183).
- Iyyer, Mohit, Peter Enns, Jordan Boyd-Graber, and Philip Resnik. 2014. "Political Ideology Detection Using Recursive Neural Networks." Pp. 1113–1122 in *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*. Baltimore, Maryland: Association for Computational Linguistics.
- James, Gareth, Daniela Witten, Trevor Hastie, and Robert Tibshirani. 2013. *An Introduction to Statistical Learning*. New York: Springer.
- Japkowicz, Nathalie, and Shaju Stephen. 2002. "The Class Imbalance Problem: A Systematic Study." *Intelligent Data Analysis* 6(5):429–49.
- Kennedy, Ryan, Stefan Wojcik, and David Lazer. 2017. "Improving Election Prediction Internationally." *Science*, 515–20.
- Khashman, Zeliha, and Adnan Khashman. 2016. "Anticipation of Political Party Voting Using Artificial Intelligence." *Procedia Computer Science* 102:611–16. doi: [10.1016/j.procs.2016.09.450](https://doi.org/10.1016/j.procs.2016.09.450).
- Khashman, Zeliha, and Adnan Khashman. 2017. "Modeling People's Anticipation for Cyprus Peace Mediation Outcome Using a Neural Model." *Procedia Computer Science* 120:734–41. doi: [10.1016/j.procs.2017.11.303](https://doi.org/10.1016/j.procs.2017.11.303).
- Kimber, Richard. 1991. "Artificial Intelligence and the Study of Democracy." *Social Science Computer Review* 9(3):381–98. doi: [10.1177/089443939100900303](https://doi.org/10.1177/089443939100900303).
- King, Gary, and Langche Zeng. 2001a. "Improving Forecasts of State Failure." *World Politics* 53(4):623–58. doi: [10.1353/wp.2001.0018](https://doi.org/10.1353/wp.2001.0018).
- King, Gary, and Langche Zeng. 2001b. "Logistic Regression in Rare Events Data." *Political Analysis* 9(2):137–63. doi: [10.1093/oxfordjournals.pan.a004868](https://doi.org/10.1093/oxfordjournals.pan.a004868).

- Knight, Will. 2017. "The Dark Secret at the Heart of AI." *MIT Technology Review*, April.
- Lagazio, Monica, and Tshilidzi Marwala. 2006. "Assessing Different Bayesian Neural Network Models for Militarized Interstate Dispute. Outcomes and Variable Influences." *Social Science Computer Review* 24(1):119–31.
- Lagazio, Monica, and Bruce Russett. 2003. "A Neural Network Analysis of Militarized Disputes, 1885-1992: Temporal Stability and Causal Complexity." in *Toward a Scientific Understanding of War: Studies in Honour of J. David Singer*. Ann Arbor: University of Michigan Press.
- Leetaru, Kalev. 2011. "Culturomics 2.0: Forecasting Large-Scale Human Behavior Using Global News Media Tone in Time and Space." *First Monday* 16(9):1–8.
- Leetaru, Kalev. 2013. "Can We Forecast Conflict? A Framework for Forecasting Global Human Societal Behavior Using Latent Narrative Indicators." University of Illinois at Urbana-Champaign.
- Linzer, Drew A. 2013. "Dynamic Bayesian Forecasting of Presidential Elections in the States." *Journal of the American Statistical Association* 108(501):124–34. doi: 10.1080/01621459.2012.737735.
- López-Iturriaga, Félix J., and Iván Pastor Sanz. 2018. "Predicting Public Corruption with Neural Networks: An Analysis of Spanish Provinces." *Social Indicators Research* 140(3):975–98. doi: 10.1007/s11205-017-1802-2.
- de Marchi, Scott, Christopher Gelpi, and Jeffrey D. Grynawski. 2004. "Untangling Neural Nets." *American Political Science Review* 98(2):371–78.
- Marwala, Tshilidzi, and Monica Lagazio. 2011. *Militarized Conflict Modeling Using Computational Intelligence*. London: Springer-Verlag.
- McCulloch, Warren S., and Walter Pitts. 1943. "A Logical Calculus of the Ideas Immanent in Nervous Activity." *The Bulletin of Mathematical Biophysics* 5(4):115–33. doi: 10.1007/BF02478259.
- Medsker, Larry R. 1994. *Hybrid Neural Network and Expert Systems*. Boston, MA: Springer.
- Michel, Jean-Baptiste, Yuan Kui Shen, Aviva Presser Aiden, Adrian Veres, Matthew K. Gray, Joseph P. Pickett, Dale Hoiberg, Dan Clancy, Peter Norvig, Jon Orwant, Steven Pinker, Martin A. Nowak, and Erez Lieberman Aiden. 2011. "Quantitative Analysis of Culture Using Millions of Digitized Books." *Science* 331(6014):176–82.
- Mitchell, Tom M. 1997. *Machine Learning*. New York: McGraw-Hill Companies.
- Monterola, Christopher, May Lim, Jerrold Garcia, and Caesar Saloma. 2002. "Feasibility of a Neural Network as Classifier of Undecided Respondents in a Public Opinion Survey." *International Journal of Public Opinion Research* 14(2):222–29. doi: 10.1093/ijpor/14.2.222.
- Morris, G. Elliott. 2019. "A Guide to Analyzing (American) Political Data in R." Retrieved March 3, 2019 (<https://bit.ly/2lBaXwL>).
- Mostafa, Mohamed M., and Ahmed A. El-Masry. 2013. "Citizens as Consumers: Profiling e-Government Services' Users in Egypt via Data Mining Techniques." *International Journal of Information Management* 33(4):627–41. doi: 10.1016/j.ijinfomgt.2013.03.007.
- Muchlinski, David, David Siroky, Jingrui He, and Matthew Kocher. 2016. "Comparing Random Forest with Logistic Regression for Predicting Class-Imbalanced Civil War Onset Data." *Political Analysis* 24(1):87–103. doi: 10.1093/pan/mpv024.
- Muchlinski, David, Xiao Yang, Sarah Birch, Craig Macdonald, and Iadh Ounis. 2020. "We Need to Go Deeper: Measuring Electoral Violence Using Convolutional Neural

- Networks and Social Media.” *Political Science Research and Methods* 1–18. doi: 10.1017/psrm.2020.32.
- Prinz, Florian, Thomas Schlange, and Khusru Asadullah. 2011. “Believe It or Not: How Much Can We Rely on Published Data on Potential Drug Targets?” *Nature Reviews Drug Discovery* 10(9):712. doi: 10.1038/nrd3439-c1.
- Raff, Edward. 2019. “A Step Toward Quantifying Independently Reproducible Machine Learning Research.” *ArXiv*.
- Ran, R. 2015. “Study on the Evaluation System of Websites of College Ideological and Political Education Based on RBF Neural Network.” *International Journal of Simulation: Systems, Science and Technology* 16(5B):8.1-8.5. doi: 10.5013/IJSSST.a.16.5B.08.
- Refenes, A. N., C. Kollias, and A. Zapranis. 1995. “External Security Determinants of Greek Military Expenditure: An Empirical Investigation Using Neural Networks.” *Defence and Peace Economics* 6(1):27–41. doi: 10.1080/10430719508404810.
- Rosenblatt, Frank. 1961. *Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms*. Washington D.C.: Spartan Books.
- Rost, Nicolas, Gerald Schneider, and Johannes Kleibl. 2009. “A Global Risk Assessment Model for Civil Wars.” *Social Science Research* 38(4):921–33. doi: 10.1016/j.ssrsearch.2009.06.007.
- Rothman, Denis. 2020. “Artificial Intelligence By Example.” *O’Reilly Media*. Retrieved November 27, 2020 (<https://www.oreilly.com/library/view/artificial-intelligence-by/9781788990547/97eeab76-9e0e-4f41-87dc-03a65c3efec3.xhtml>).
- Ruder, Sebastian. 2017. “Transfer Learning - Machine Learning’s Next Frontier.” Retrieved February 9, 2020 (<https://ruder.io/transfer-learning/>).
- Schneider, Gerald, Nils Petter Gleditsch, and Sabine Carey. 2011. “Forecasting in International Relations: One Quest, Three Approaches.” *Conflict Management and Peace Science* 28(1):5–14. doi: 10.1177/0738894210388079.
- Schrodtt, Philip A. 1990. “Predicting Interstate Conflict Outcomes Using a Bootstrapped ID3 Algorithm.” *Political Analysis* 2:31–56. doi: 10.1093/pan/2.1.31.
- Schrodtt, Philip A. 1991. “Prediction of Interstate Conflict Outcomes Using a Neural Network.” *Social Science Computer Review* 9(3):359–80.
- Schrodtt, Philip A. 2000. “Pattern Recognition of International Crises Using Hidden Markov Models.” Pp. 296–328 in *Political complexity: nonlinear models of politics*, edited by D. Richards. Ann Arbor: University of Michigan Press.
- Schrodtt, Philip A. 2004. *Patterns, Rules and Learning: Computational Models of International Behavior*. 2nd Ed. Vinland, Kansas: Parus Analytical Systems.
- Spiegeleire, Stephan De. 2016. *Great Power Assertivitis*. The Hague.
- Srivastava, Nitish, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. “Dropout: A Simple Way to Prevent Neural Networks from Overfitting.” *Journal of Machine Learning Research* 15:1929–58.
- Ulfelder, Jay. 2012. “Why the World Can’t Have a Nate Silver.” *Foreign Policy*.
- Vanhanen, Tatu, and Richard Kimber. 1994. “Predicting and Explaining Democratization in Eastern Europe.” Pp. 63–96 in *Democratization in Eastern Europe: domestic and international perspectives*, edited by G. Pridham and T. Vanhanen. London; New York: Routledge.
- Wang, Wei, David Rothschild, Sharad Goel, and Andrew Gelman. 2015. “Forecasting Elections with Non-Representative Polls.” *International Journal of Forecasting* 31(3):980–91. doi: 10.1016/j.ijforecast.2014.06.001.

- Ward, Michael D., and Andreas Beger. 2017. "Lessons from near Real-Time Forecasting of Irregular Leadership Changes." *Journal of Peace Research* 54(2):141–56. doi: 10.1177/0022343316680858.
- Ward, Michael D., Brian D. Greenhill, and Kristin M. Bakke. 2010. "The Perils of Policy by P-Value: Predicting Civil Conflicts." *Journal of Peace Research* 47(4):363–75. doi: 10.1177/0022343309356491.
- Weber, Patrick, Nicolas Weber, Michael Goesele, and Rüdiger Kabst. 2018. "Prospect for Knowledge in Survey Data." *Social Science Computer Review* 36(5):575–90. doi: 10.1177/0894439317725836.
- Weidmann, Nils B., and Michael D. Ward. 2010. "Predicting Conflict in Space and Time." *Journal of Conflict Resolution* 54(6):883–901.
- Werbos, Paul. 1974. "Beyond Regression: New Tools for Prediction and Analysis in the Behavioral Sciences." Harvard University.
- Werbos, Paul. 1994. *The Roots of Backpropagation: From Ordered Derivatives to Neural Networks and Political Forecasting*. New York: Wiley.
- Witmer, Frank D. W., Andrew M. Linke, John O'Loughlin, Andrew Gettelman, and Arlene Laing. 2017. "Subnational Violent Conflict Forecasts for Sub-Saharan Africa, 2015–65, Using Climate-Sensitive Models." *Journal of Peace Research* 54(2):175–92. doi: 10.1177/0022343316682064.
- Zeng, Langche. 1999. "Prediction and Classification with Neural Network Models." *Sociological Methods & Research* 27(4):499–524. doi: 10.1177/0049124199027004002.
- Zeng, Langche. 2000. "Neural Network Models for Political Analysis." Pp. 239–68 in *Political complexity: nonlinear models of politics*, edited by D. Richards. Ann Arbor: University of Michigan Press.
- Zhang, Han, and Jennifer Pan. 2019. "CASM: A Deep-Learning Approach for Identifying Collective Action Events with Text and Image Data from Social Media." *Sociological Methodology* 49(1):1–57. doi: 10.1177/0081175019860244.
- Zhao, C., Y. Li, and Z. Li. 2015. "Study on the Evaluation Model of Undergraduates' Ideological and Political Status Based on BP Neural Network." *International Journal of Simulation: Systems, Science and Technology* 16(5B):7.1-7.8. doi: 10.5013/IJSSST.a.16.5B.07.