

Determining Crucial Factors for the Popularity of Scientific Articles

R. JANKOWSKI AND J. SIENKIEWICZ*

Faculty of Physics, Warsaw University of Technology, Koszykowa 75, PL 00662 Warszawa, Poland

Doi: [10.12693/APhysPolA.138.41](https://doi.org/10.12693/APhysPolA.138.41)

*e-mail: julian.sienkiewicz@pw.edu.pl

Using a set of over 70.000 records from PLOS One journal consisting of 37 lexical, sentiment and bibliographic variables we perform analysis backed with machine learning methods to predict the class of popularity of scientific papers defined by the number of times they have been viewed. Our study shows correlations among the features and recovers a threshold for the number of views that results in the best prediction outcomes in terms of Matthew's correlation coefficient. Moreover, by creating a variable importance plot for random forest classifier, we are able to reduce the number of features while keeping similar predictability and determine crucial factors responsible for the popularity.

topics: scientometrics, scientific data, predictions, machine learning

1. Introduction

The single most popular bibliometric criterion for judging the impact of scientific papers is the number of received citations, commonly known as “citation count”. However, due to a rather indirect discipline dependence [1], this metric alone can be in many cases unreliable. Many services have created their own metrics to determine the most popular article which shows a complex landscape of measures [2]. Google Scholar (scholar.google.com) puts weight especially on citation count and words included in a document's title. As a consequence, the first search results are often highly cited articles. On the other hand, the Public Library of Science (www.plos.org) takes into account the number of HTML page views and PDF downloads. In Scopus (www.scopus.com), there are a few different article-level metrics: (i) Scholarly Activity (Mendeley readers), (ii) Scholarly Commentary (Blogs, Wikipedia), and (iii) Mass Media (Twitter, Facebook, Reddit), which brings this service closer to so-called “altmetrics” [3] which has lately become one of the key components used in evaluating scientific papers.

There is an overwhelming motion among scientists [4–6] and in particular in the popular media coverages to connect the popularity of articles to some of their textual features in a simple way. A running example is to directly state that shorter titles increase the chance of a scientific article to be cited [7]. A counterexample to this simplified approach is a recent paper [8], where authors performed an investigation of how different textual properties of scientific papers affect the number of citations they acquire. Using a set of

over 4.3 million and applying quantile regression in order to examine different regimes of citations it has been shown that in most journals, short titles correlate negatively with citations only for the most cited papers. On the other hand, for typical papers, this relation is usually absent. It has also been noticed that depending on the journal the results may vary, which emphasizes the nonlinearity of the relationship.

In this work, we approach the problem from a different point of view. The main distinction is that we not only intend to explore the popularity aspects but also use tools that would, in turn, allow predicting whether the paper would be popular based on available features. In addition to that, as considering just a single feature (e.g., the length of the title) at a time might not be sufficient to cover complex relations among different features, we decided to use the machine learning (ML) approach. Moreover, this study takes into account a single journal (i.e., PLOS One), as compared to the previous investigation [8] in which authors used over 1500 different journals. Last but not least, we recorded the information about the number of views per month, in contrast to the sum of citations the article acquired. This distinction seems to be important as almost everyone has a chance to view the page with the given article, but only people who are experts in a specific topic can cite it. Thus, the number of views is more a popularity measure.

The rest of the paper is organized as follows: first we describe in detail the data and show some exploratory data analysis to examine correlations among the features and distribution of views, as

well as its consequences. Then we introduce machine learning methods and two basic metrics that allow for quantification of predictability of tested methods. The fourth chapter contains the results of this study: the quality of predictions on the full dataset using selected classifiers and similar outcomes for reduced data. We end the paper with some concluding remarks.

2. Data and exploratory analysis

In this study we use two datasets: the first one has been downloaded using the Public Library of Science automatic services (PLOS API) (api.plos.org), and it included the following information about each paper published in the PLOS ONE journal: (i) title, (ii) list of authors, (iii) abstract contents, (iv) full text contents. The downloaded data comprises around 140.000 records of data. The second set was acquired using Article-Level Metrics data (PLOS ALM) (alm.plos.org) — a service that keeps monthly statistics of views (i.e., visiting web-page with the given paper). After data cleaning and matching, the dataset contained over 70,000 papers published between years 2003 and 2014. To overcome issues with variable impact, i.e., some papers can initially obtain large attention which then rapidly decays while other can be characterized by a stable

pattern of number of views, we decided to use the average number of views per month.

Table I gathers all the features calculated for each item in the dataset (37 features in total). They are connected to different aspects of the scientific text: number of words, characters and sentences are purely length properties. Another widely utilized property of the text is its emotional content (sometimes called sentiment). Here we use valence (i.e., the emotional sign of the text) and arousal which indicates the level of activation evoked by the text. These parameters have been used to quantify the collective behavior of online users [9], to model the evolution of online discussions in certain environments [10] or to predict the dynamics of Twitter users during the Olympic Games in London [11]. In order to find the sentiment of the title and abstract the study by Warriner et al. was used [12] where authors created a database containing nearly 14,000 English lemmas with valence and arousal values. Finally, the last property taken into account in this study are the references appearing in the text. Here, we have been interested both in the plain number of citations, the part of the manuscript where they appeared (introduction, main part, discussion), as well as in their emotional context, i.e., the sentiment in the proximity of the given reference. This kind of analysis can help to judge the role of negative citations [13].

TABLE I

Features calculated for each article. Numbers in parentheses in the second column give the number of features coming from this property.

	Property	Comments
1	number of words in the abstract, title and the full text (3)	—
2	number of characters in the abstract, title and the full text (3)	—
3	number of sentences in the abstract and the full text (2)	—
4	number of authors (1)	—
5	number of citation in the full text (1)	To calculate the number of citation the regular expression was used: [number]
6	number of citations in the introduction and discussion (2)	See above comment
7	valence and arousal for the title, abstract and full text (6)	Using dictionary provided by Warriner et al. [12]
8	valence and arousal related to specific citation in the full text (2)	The sentiment was calculated taking into account 100 characters before and after the citation.
9	valence and arousal related to specific citation after splitting full text into four parts (8)	The full text was divided into 4 parts; for each part positions of the citations were identified and then the sentiment was calculated as above.
10	valence and arousal after splitting full text into four parts (8)	—
11	average number of views (per month) from the date of publication and the end of year 2014 (1)	—




Fig. 1. Correlation matrix of all features. To obtain better visual effect small correlation values ($r \in [-0.1, 0.1]$) were set to zero.

Figure 1 shows a correlation matrix based on all the features recognized in text: each cell of the matrix is simply calculated as a Pearson correlation coefficient of two features X_i and X_j

$$\rho_{X_i, X_j} = \frac{\text{cov}(X_i, X_j)}{\sigma_{X_i} \sigma_{X_j}}, \quad (1)$$

where $\text{cov}(X_i, X_j)$ is the covariance of features X_i and X_j and σ_{X_i} , σ_{X_j} are, respectively, their variances. This shows an obvious, strong correlation between the number of words, characters, and sentences in the abstract, title and full text (middle part of the plot). Moreover, there is a noticeable relationship linking the number of citations and the number of words/characters in the full text. Valence features are correlated positively with each other (bottom-right part of the plot), as well as arousal features (upper-left part of the plot), however, the correlation between the valence and arousal features is slightly negative (upper-right part of the plot). Interestingly, features connected with sentiment are weakly or not correlated with length properties. It is also notable that none of these features has any connections to the main observable of this study, i.e., the average number of views.




Fig. 2. Probability distribution $p(v)$ of the number of views v (log-log scale used). The median value (marked as a dashed line) is around 79 views.

The reason for the above-mentioned fact comes directly from the character of distribution of views depicted in Fig. 2: it is a heavy-tailed function, however, unlike popular power-law distributions, often encountered in sociophysics, econophysics and science of science [14], it possesses a peak located close to the median value. It immediately gives rise to a crucial question: which papers should be treated

as “popular” ones taking into account the number of views as a metric? Pursuing this line of thought even further, one might wonder which threshold v_{th} that divides the papers into two classes: popular (say “positive”) and not popular (say “negative”) can lead to a good prediction of such classes based on the defined features.

3. Methods

Data Mining (DM) or Machine Learning (ML) [15, 16] can be, very generally, described as a rather multidisciplinary field that shows how to learn from data, and make predictions about them. Its connections to physical sciences, which are bi-directional as both disciplines are mutually gaining due to co-operation have lately been emphasized in a recent review [17]. The major advantage of DM is that its methods can simultaneously search relations among model variables (features) and the modeled outcome regardless of the size of the space features.

One of the main categories of ML is supervised learning which concerns learning from a set of previously labeled data. In our case, as described in Sect. 2, we would like to be able to predict the class of the paper (popular/not popular), which is a typical example of binary classification. In such a setting, a confusion matrix is a 2 by 2 table (see Table II), that allows visualization of the performance of an algorithm, basing on the number of occurrences when originally positive / negative case was truly (or falsely) predicted as positive / negative.

However, it is usually much better to work on a single index instead of a compound metrics such as a matrix. In this work two such metrics were used: F_1 score [18] and Matthews correlation coefficient (MCC) [19] defined as follows

$$F_1 = \frac{2TP}{2TP + FP + FN} \quad (2)$$

$$MCC = \frac{TP \times TN - FP \times FN}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}} \quad (3)$$

with $F_1 \in [0, 1]$ and $MCC \in [-1, 1]$, and where F_1 score is a harmonic mean of the precision (i.e., $\frac{TP}{TP+FP}$) and recall (i.e., $\frac{TP}{TP+FN}$). On the other hand, MCC is more informative than other confusion matrix measures because it takes into account the balance ratios of the four confusion matrix categories. The value of MCC equal to zero means that the model in question chooses the class randomly. The closer value to 1, the better the model performs.

The whole dataset was randomly split into a training and testing part in proportion 75% to 25%. Based on Fig. 2 the popularity threshold for ML models was chosen (20–300 views) which covers a sufficiently large range of divisions between popular and not popular papers (see Fig. 3).




Fig. 3. The percentage of popular articles in the dataset depending on the popularity threshold. Dashed lines mark the range of 70–90 of views and the corresponding ratios of popular articles.

Confusion matrix.

TABLE II

	actual positive	actual negative
predicted positive	true positive (TP)	false positive (FP)
predicted negative	false negative (FN)	true negative (TN)

For each such value, a set of binary classifiers (see below) was created and their performance was assessed on the testing set.

To compare models with different complexity we chose the following classification algorithms: (i) Linear Discriminant Analysis (LDA) [20, 21] projects high-dimensional data onto a low-dimensional space where the data achieves maximum class separability, (ii) Quadratic Discriminant Analysis (QDA), relatively similar to LDA, however, does not assume the covariance of each of the classes to be identical, (iii) Support Vector Machines (SVM) [22] map the data into a higher dimensional input space and construct an optimal separating hyperplane in this space. Due to the kernel method SVM allow for creating nonlinear decision surfaces, (iv) Logistic regression [23] is a mathematical modeling approach that can be used to describe the relationship of several features to a dichotomous dependent variable, such as if an article is popular or not, (v) Random forest classifier [24] consists of a combination of tree classifiers where each classifier is generated using a random vector sampled independently from the input vector (features), and each tree casts a vote for the most popular class to classify an input vector.

4. Results

Figure 4 shows the results of each model’s MCC and F_1 values, respectively, as a function of the views threshold for a relevant range of values. It is easy to notice that F_1 (Fig. 4b) brings hardly

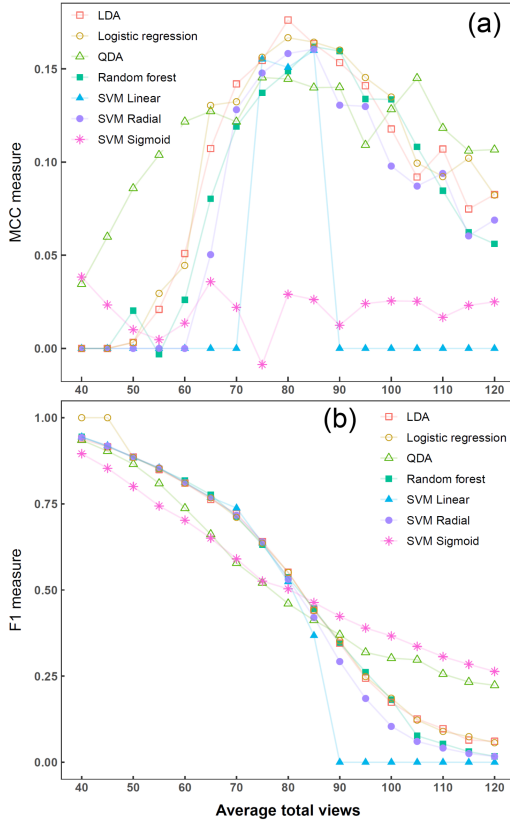


Fig. 4. (a) MCC value versus the threshold of average number of views. (b) F_1 score versus the threshold of average number of views. Different symbols reflect different ML algorithms.

any insights as it is basically monotonically decreasing function with the threshold value. This effect is probably caused by an unbalanced dataset — as was mentioned before, F_1 takes into account only TP, FP and FN values without the input from TN. Thus, for further analysis, we disregard the outcomes of F_1 . On the other hand, MCC (Fig. 4a) is much more informative: it is clear that the majority of used ML models (except for the sigmoidal version of SVM) tend to have a maximum between 70 and 90 views which roughly corresponds to 40%–60% popularity (see Fig. 3). Although the models come from different domains the differences among them are rather minimal with LDA being the best one. One needs to underline that the maximal value of $MCC \approx 0.18$ suggests a rather mild predictive power of the approach. Almost all the methods point to $v \approx 80$ as the best threshold value, thus suggesting a 50 ÷ 50 division between popular and not popular papers. QDA seems to outperform other methods in case of searching for the best model for different regimes of v .

Apart from just acquiring the best possible prediction rate, one is usually also interested in understanding primary drivers for such behavior. If the model performs in the same or nearly the same way while some of its features are removed

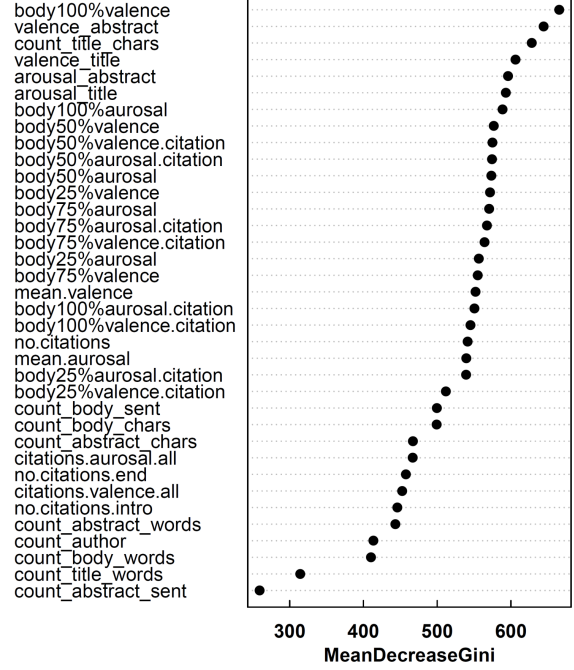


Fig. 5. The value of Mean Decrease Gini of each feature. The random forest model was built for 80 views.

one can conclude that the remaining features are crucial, i.e., they are key variables deciding on the modeled outcome. Such an approach is generally referred as to dimension reduction — its probably most known example is the Principal Component Analysis (PCA, described in detail, e.g. in [25]) that allows to remove highly correlated variables and with respect to the science of science can be used to show distinctions among scientific areas or fields [26]. However, as one of the methods used in this study is the random forest, we are able to employ the Mean Decrease Gini instead. Gini Index [27] is an attribute selection measure, which measures the impurity of an attribute with respect to the classes. This metric allows to distinguish most influential features for random forest classifier and is often used to reduce the number of features in the input vector (cf. [28]).

We selected the random forest model for the best MCC value (80 views). Figure 5 shows the Mean Decrease Gini for each feature: the first three features (i.e., valence in the 4th part of the full text, valence in the abstract as well as the number of characters in the title) were taken to build once again selected machine learning models. A comparison of the models constructed on reduced features and full features is shown in Fig. 6. For the majority of the models, dimension reduction leads to a drastic decrease of MCC measure, thus implying that the reduced dataset has no predictive power. However, surprisingly in the case of logistic-regression (Fig. 6f) both full and reduced sets acquire almost the same MCC value.

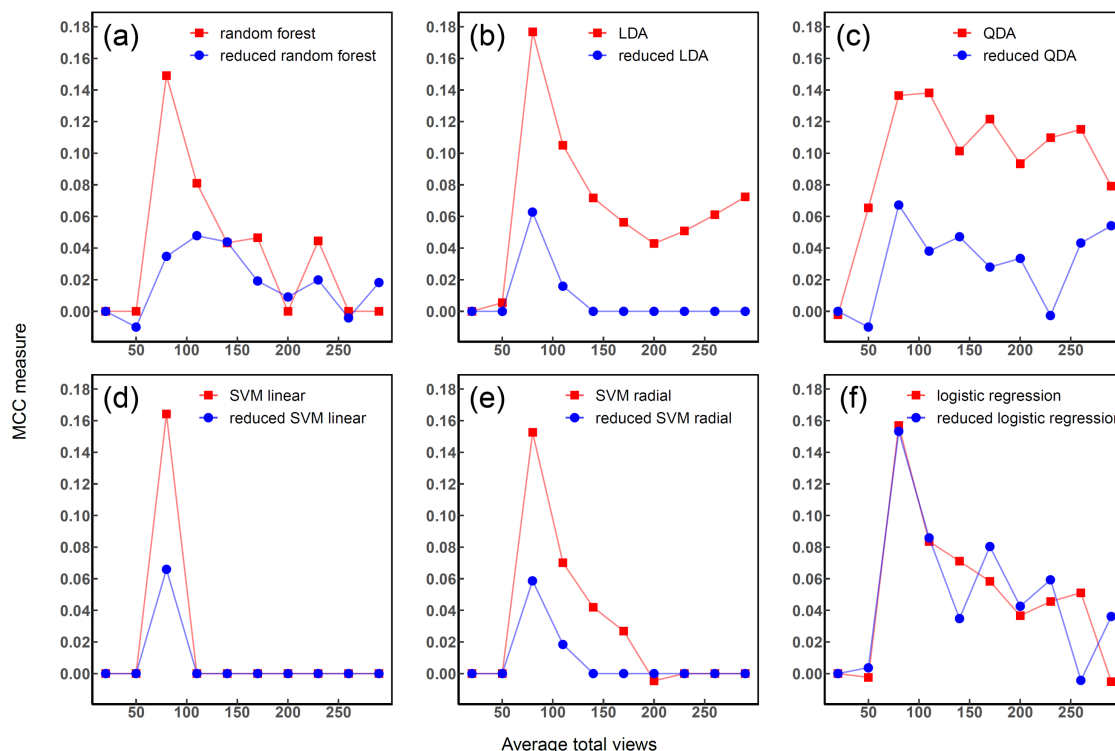


Fig. 6. Comparison of MCC values for models built on all features (depicted by squares) and the ones selected by dimension reduction (given by circles): (a) random forest, (b) LDA, (c) QDA, (d) linear SVM, (e) radial SVM and (f) logistic regression.

5. Conclusions

In this study we have conducted an analysis of over 70.000 articles coming from PLOS One journal, trying to link their lexical, sentiment and bibliographic properties with the popularity they acquired, measured as the number of views of a single paper. In contrast to previous works we have put the emphasis on the question of predicting the popularity and testing the quality of such methods, which has led us to machine learning approaches. In order to overcome the problem of ambiguity of the popularity we have decided to use classification approach by introducing a threshold that divides the number of views and, consequently, the papers themselves into two classes: popular and not popular ones. Basing on the selection of machine learning models the popularity threshold was set between 70 and 90 views with the maximum value close to 80, being the median of the views distribution. For such a threshold the majority of the examined models achieved the best predictive power. Using random forest model we have selected three most crucial features, e.g., valence in the 4th part of full text and in abstract as well as the number of characters in the title. Models built on a reduced number of features performed in general worse than on the total number of features with the prominent exception of logistic regression.

Let us underline that these results do not contradict the conclusions shown in [8] where the most rel-

evant identified factors were number of authors and length of abstract (in Fig. 5 counted as almost least important ones). The difference comes from the fact that in our study we considered a single journal whereas [8] aggregated results from over 1.500 journals (not even including PLOS One). On the other hand, our study also brings the length of the title to the front. It is also essential to point out that we paid attention to a different measure of attention, i.e., the number of views that counts how many times a specific paper has been viewed. Such a variable, as opposed to the number of citations might be seen as less reliable — entering the web page containing the paper does not automatically result in reading it. The difference between the views and citations can also be more fundamental: making a reference to specialized work suggests expertise in the given scientific area whereas such knowledge is not needed for merely entering a paper, e.g., attracted by the title. However, these differences do not indicate that there exists a proper measure of attention for scientific papers, instead marking that they both follow specific dynamics.

In our opinion the study can be easily extended by examining the results of the Principal Components Analysis or comparing the outcomes of quantile regression method. Moreover, as the service of the Public Library of Science gives access also to the number of PDF downloads it would be interesting to combine these measures together with the number of citations.

References

- [1] F. Didegah, M. Thelwall, *J. Informetr.* **7**, 861 (2013).
- [2] J. Bollen, H. Van de Sompel, A. Hagberg, R. Chute, *PLOS ONE* **4**, 111 (2009).
- [3] W. Glänzel, J. Gorraiz, *Scientometrics* **102**(3), 21612164 (2015).
- [4] C. Paiva, J. Lima, B. Paiva, *Clinics* **67**, 509513 (2012).
- [5] A. Letchford, H.S. Moat, T. Preis, *R. Soc. Open Sci.* **2**(8), 150266 (2015).
- [6] F. Rostami, A. Mohammadpoorasl, M. Hajizadeh, *Scientometrics* **98** 2007 (2014).
- [7] B. Deng, “Papers with shorter titles get more citations”, *Nature. News*, 2015.
- [8] J. Sienkiewicz, E.G. Altmann, *R. Soc. Open Sci.* **3**, 160140 (2016).
- [9] A. Chmiel, J. Sienkiewicz, M. Thelwall, G. Paltoglou, K. Buckley, A. Kappas, J.A. Holyst, *PLOS ONE* **6**, 18 (2011).
- [10] J. Sienkiewicz, M. Skowron, G. Paltoglou, J.A. Holyst, *Adv. Complex Sys.* **16**, 1350026 (2013).
- [11] J. Chołoniewski, J. Sienkiewicz, J.A. Holyst, M. Thelwall, *Acta Phys. Pol. A* **127A**, 21 (2014).
- [12] A.B. Warriner, V. Kuperman, M. Brysbaert, *Behav. Res. Methods* **45**, 11911207 (2013).
- [13] Ch. Catalini, N. Lacetera, A. Oettl, *Proc. Natl. Acad. Sci. (PNAS)* **112**, 13823 (2015).
- [14] M.E.J. Newman, *Contemp. Phys.* **46** 323 (2005).
- [15] T. Hastie, R. Tibshirani, J. Friedman, *The Elements of Statistical Learning*, Springer Series in Statistics, Springer, New York 2001.
- [16] P. Mehta, M. Bukov, Ch.-H. Wang, A.G.R. Day, C. Richardson, Ch.K. Fisher, D.J. Schwab, *Phys. Rep.* **810**, 1124 (2019).
- [17] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, L. Zdeborová, *Rev. Mod. Phys.* **91**, 045002 (2019).
- [18] C.J. Van Rijsbergen, *Information Retrieval* (2nd ed.) Butterworth-Heinemann, London, 1979.
- [19] B.W. Matthews, *Biochim. Biophys. Acta, Protein Struct.* **405**, 442 (1975).
- [20] S. Mika, G. Ratsch, J. Weston, B. Scholkopf, K.-R. Mullers, in: *Neural Networks for Signal Processing IX: Proc. of the 1999 IEEE Signal Processing Society Workshop, Madison (WI)*, IEEE, 1999, p. 41.
- [21] J. Ye, in: *ICML '07: Proceedings of the 24th international conference on Machine learning ACM*, Association for Computing Machinery, New York 2007 p. 1087.
- [22] M.A. Hearst, S. T. Dumais, E. Osuna, J. Platt, B. Scholkopf, *IEEE Intell. Sys. App.* **13**, 18 (1998).
- [23] D.G. Kleinbaum, M. Klein, *Logistic regression*, Springer, New York 2002.
- [24] L. Breiman, *Mach. Learn.* **45**, 532 (2001).
- [25] L.G. Gajewski, J. Chołoniewski, J.A. Holyst, *Acta Phys. Pol. A* **129**, 1071 (2016).
- [26] J. Sienkiewicz, K. Soja, J.A. Holyst, P.M.A. Slood, *Sci. Rep.* **8**, 8253 (2018).
- [27] L. Breiman, *Mach. Learn.* **24**, 41 (1996).
- [28] J. Arumugam, S.T.S. Bukkapatnam, K.R. Narayanan, A.R. Srinivasa, *PLOS ONE* **11**, 112 (2016).