

Research Reproducibility as a Survival Analysis

Edward Raff

Booz Allen Hamilton
University of Maryland, Baltimore County
raff.edward@bah.com
raff.edward@umbc.edu

Abstract

There has been increasing concern within the machine learning community that we are in a reproducibility crisis. As many have begun to work on this problem, all work we are aware of treat the issue of reproducibility as an intrinsic binary property: a paper *is* or *is not* reproducible. Instead, we consider modeling the reproducibility of a paper as a survival analysis problem. We argue that this perspective represents a more accurate model of the underlying meta-science question of reproducible research, and we show how a survival analysis allows us to draw new insights that better explain prior longitudinal data. The data and code can be found at <https://github.com/EdwardRaff/Research-Reproducibility-Survival-Analysis>

1 Introduction

There is current concern that we are in the midst of a reproducibility crisis within the fields of Artificial Intelligence (AI) and Machine Learning (ML) (Hutson 2018). Rightfully, the AI/ML community has done research to understand and mitigate this issue. Most recently, Raff (2019) performed a longitudinal study attempting to independently reproduce 255 papers, and they provided features and public data to begin answering questions about the factors of reproducibility in a quantifiable manner. Their work, like all others of which we are aware, evaluated reproducibility using a binary measure. Instead, we argue for and demonstrate the value of using a survival analysis. In this case, we model the hazard function $\lambda(t|\mathbf{x})$, which predicts the likelihood of an event (i.e., reproduction) occurring at time t , given features about the paper \mathbf{x} . In survival analysis, we want to model the likely time until the occurrence of an event. This kind of analysis is common in medicine where we want to understand what factors will prolong a life (i.e., increase survival), and what factors may lead to an early death.

In our case, the event we seek to model is the successful reproduction of a paper’s claims by a reproducer that is independent of the original paper’s authors. Compared to the normal terminology used, we desire factors that *decrease survival* meaning a paper takes *less time* to reproduce. In the converse situation, a patient that “lives forever” would be equivalent to a paper that cannot be reproduced with any

amount of effort. We will refer to this as *reproduction time* in order to reduce our use of standard terminology that has the opposite connotation of what we desire. Our goal, as a community, is to reduce the reproduction time toward zero and to understand what factors increase or decrease reproduction time.

We will start with a review of related work on reproducibility, specifically in machine learning, in §2. Next, we will detail how we developed an extended dataset with paper survival times in §3. We will use models with the Cox proportional hazards assumption to perform our survival analysis, which will begin with a linear model in §4. The linear model will afford us easier interpretation and statistical tools to verify that the Cox proportional hazards model is reasonable for our data. In §5 we will train a non-linear Cox model that is a better fit for our data in order to perform a more thorough analysis of our data under the Cox model. Specifically, we show in detail how the Cox model allows us to better explain observations originally noted in (Raff 2019). This allows us to measure a meaningful effect size for each feature’s impact on reproducibility, and better study their relative importance’s. We stress that the data used in this study does not mean these results are definitive, but useful as a new means to think about and study reproducibility.

2 Related Work

The meta-science question of reproducibility — scientifically studying the research process itself — is necessary to improve how research is done in a grounded and objective manner (Ioannidis 2018). Significant rates of non-reproduction have been reported in many fields, including a 6% rate in clinical trials for drug discovery (Prinz, Schlange, and Asadullah 2011), making this an issue of greater concern in the past few decades. Yet, as far as we are aware, all prior works in machine learning and other disciplines treat reproducibility as a binary problem (Gundersen and Kjensmo 2018; Glenn and P.A. 2015; Ioannidis 2017; Wicherts et al. 2006). Even works that analyze the difference in effect size between publication and reproduction still view the issue result in a binary manner (Collaboration 2015).

Some have proposed varying protocols and processes that authors may follow to increase the reproducibility of their work (Barba 2019; Gebru et al. 2018). While valuable, we lack quantification of their effectiveness. In fact, little has

been done to empirically study many of the factors related to reproducible research, with most work being based on a subjective analysis. Olorisade, Brereton, and Andras (2018) performed a small scale study over 6 papers in the specific sub-domain of text mining. Bouthillier, Laurent, and Vincent (2019) showed how replication (e.g., with docker containers), can lead to issues if the initial experiments use fixed seeds, which has been a focus of other work (Forde et al. 2018). The largest empirical study was done by Raff (2019), which documented features while attempting to reproduce 255 papers. This study is what we build upon for this work.

Sculley et al. (2018) have noted a need for greater rigor in the design and analysis of new algorithms. They note that the current focus on empirical improvements and structural incentives may delay or slow true progress, despite appearing to improve monotonically on benchmark datasets. This was highlighted in recent work by Dacrema, Cremonesi, and Jannach (2019) who attempted to reproduce 18 papers from Neural Recommendation algorithms. Their study found that only 7 could be reproduced with “reasonable effort.” Also concerning is that 6 of these 7 could be outperformed by better tuning baseline comparison algorithms. These issues regarding the nature of progress, and what is actually learned, are of extreme importance. We will discuss these issues as they relate to our work and what is implied about them from our results. However, the data from (Raff 2019) that we use does not quantify such “delayed progress,” but only the reproduction of what is stated in the paper. Thus, in our analysis, we would not necessarily be able to discern the issues with the 6 papers with insufficient baselines by (Dacrema, Cremonesi, and Jannach 2019).

We note that the notion of time impacting reproduction by the *original authors* or using *original code* has been previously noted (Mesnard and Barba 2017; Gronenschild et al. 2012) and often termed “technical debt” (Sculley et al. 2015). While related, this is a fundamentally different concern. Our study is over independently attempted implementation, meaning technical debt of an existing code base can not exist. Further, these prior works still treat replication as a binary question despite noting the impact of time on difficulty. Our distinction is using the time to implement itself as a method of quantifying the degree or difficulty of replication, which provides a meaningful effect size to better study replication.

3 Study Data

The original data used by Raff (2019) was made public but with explicit paper titles removed. We have augmented this data in order to perform this study. Specifically, of the papers that were reproduced, the majority had their implementations made available as a part of the JSAT library (Raff 2017). Using the Github history, we were able to determine end dates for the completion of an algorithm’s implementation. In addition the original start dates of each implementation are recorded by Mendeley which were used for the original study. Combined, this gives us start and end dates and survival times for 90 out of the 162 papers that were reproduced. The remaining 44% of reproduced papers, for which we could not determine any survival time, were unfortunately excluded from the analysis conducted in the remainder of this paper.




Figure 1: Histogram of the time taken to reproduce. The dark blue line shows a Kernel Density Estimate of the density, and the dashes on the x-axis indicate specific values.

Summary statistics on the time taken to reproduce these 90 papers are shown in Figure 1. While most were reproduced quickly, many papers required months or years to reproduce.

Raff (2019) noted that the attempts at implementation were not continuous efforts, and noted many cautions of potential bias in results (e.g., most prominently all attempts are from a single author). Since we are extending their data, all these same biases apply to this work — with additional potential confounders. Total amount of time to implement could be impacted by life events (jobs, stresses, etc.), lack of appropriate resources, or attempting multiple works simultaneously, which is all information we lack. As such readers must temper any expectation of our results being a definitive statement on the nature of reproduction, but instead treat results as initial evidence and indicators. Better data must be collected before stronger statements can be made. Since the data we have extended took 8 years of effort, we expect larger cohort studies to take several years of effort. We hope our paper will encourage such studies to include time spent into the study’s design and motivate participation and compliance of cohorts.

With appropriate caution, our use of Github as a proxy measure of time spent gives us a level of ground truth for the reproduction time for a subset of reproduced papers. We lack any labels for the time spent on papers which failed to be reproduced and successful reproductions outside of Github. Survival analysis allows us to work around some of these issues as a case of *right censored* data. Right censored data would indicate a start time s , an amount of time observed t_o , but not the successful event (reproduction) e . If t_e is the amount of time needed to observe the event e (i.e., true amount of time needed to reproduce), then we have $t_o < t_e$. If we can estimate a minimum time $0 < \hat{t}_o < t_o$, we can still perform a reasonable and meaningful survival analysis. We work under the assumption that more time was likely spent on papers that were not reproduced than ones that were, and so assume that the average time spent on successful papers (\hat{t}_o) than the actual amount of time spent (t_o). As such, we assign every non-reproduced paper the average amount of

time for the data in Figure 1, or 218.8 days of effort, as an approximate guess at the amount of time that would be expended. We note that we found our analysis to be robust to a wide range of alternative constants, such as the median time (53.5 days). While the exact values in the analysis would change, the results were qualitatively the same. A repetition showing qualitative similarity of using the median can be found in the appendix.

A survival model does not provide us the means to circumvent the cases where we lack the observed time t_o for successfully reproduced papers. Due to the reduced number of data points (the 72 papers not implemented in JSAT’s Github), we have excluded some of the variables from analysis in this paper. Specifically, the Venue a paper was published in and the Type of Venue. This reduced dataset resulted in significant skew in some of the sub-categories of these fields, making comparisons untenable (e.g., the number of papers in Workshops was reduced to one example).

The above protocol was used to create the data for this study, and will be made available for others to study and perform additional analysis on. Our lack of ground truth “level of effort” for all of the original data is a source of potential bias in our results, and should be used to caution on taking the results as any kind of proclamation or absolute truth. That said, we find the analysis useful and able to elucidate finer details and insights that were not recognizable under the standard binary analysis of reproduction.

4 Linear Hazard Analysis

We start with standard Cox model, where we use a linear set of coefficients $\beta \in \mathbb{R}^d$ to control how much our various features impact the predicted reproduction time, $\lambda(t|x_i) = \exp(\mathbf{x}_i^\top \beta) \cdot \lambda_0(t)$, where $\lambda_0(t)$ is a baseline hazard function. The Cox model imposes a *proportional* assumption on the nature of the hazard. The base hazard rate $\lambda_0(t)$ may be different for every individual, but the proportional assumption means that we expect the impact of altering any covariate to have the same proportional effect for every instance. e.g., specifying the hyperparameters used in a paper would always reduce the reproduction time by a factor of X%.

The Cox model provides us a number of benefits to our analysis, provided that it is a reasonable assumption. First, it allows us to estimate β without knowing, or even modeling, the base hazard function $\lambda_0(t)$. Second, though not unique to the Cox model, it supports *right censored* data. If an instance is right censored, we have waited for the event to occur for some unit of time t , but have not yet seen the event occur at a later point in time t_e . This allows us to model all of the failed reproduction attempts as papers which may be reproducible, but for which we have not yet put in sufficient effort to reproduce the paper. This allows us to use our estimated effort spent on non-reproduced papers without causing significant harm to the underlying model.

4.1 Cox Proportional Hazard Validation

The question we must first answer: is the Cox proportional model assumption reasonable for our data? First, as a baseline, we have the results of the non-parametric tests per-

Feature	Independent	Logistic	Cox
Year Published	0.964	0.613	0.92
Year Attempted	0.674	0.883	0.45
Has Appendix	0.330	0.201	0.07
Uses Exemplar Toy Problem	0.720	0.858	0.20
Looks Intimidating	0.829	0.035	0.20
<i>Exact Compute Used</i>	0.257	1.000	0.39
<i>Data Available</i>	< 0.005	0.644	0.81
Code Available	0.213	0.136	0.18
Number of Authors	0.497	0.542	0.68
Pages	0.364	0.702	0.82
Num References	0.740	0.421	0.54
Number of Equations	0.004	0.478	0.74
Number of Proofs	0.130	0.629	0.47
Number of Tables	0.010	0.618	0.86
Number of Graphs/Plots	0.139	0.325	0.51
Number of Other Figures	0.217	0.809	0.98
Conceptualization Figures	0.365	0.349	0.13
Hyperparameters Specified	< 0.005	1.000	0.49
Algorithm Difficulty	< 0.005	1.000	0.03
Paper Readability	< 0.005	1.000	< 0.005
Pseudo Code	< 0.005	1.000	< 0.005
Compute Needed	< 0.005	1.000	1.000
Rigor vs Empirical	< 0.005	1.000	0.02

Table 1: p-values for each features’ importance.

formed by (Raff 2019). Next, we train a linear logistic regression model and a linear Cox proportional hazards model on our reduced subset, and we compute and compare which features were found as significant. Due to the impact of outliers in our data, we found it necessary to use a robust estimate of the Cox linear model (Lin and Wei 1989). Last, we perform simple cross-validation to show reasonable predictive performance.

We highlight here that our goal is not to use these results as a new determination for what factors are important, but as a method of comparing the relative appropriateness and agreement between a logistic and Cox model with that of the original analysis of (Raff 2019). The two new models are trained on less data. In addition, the logistic and Cox models required a regularization penalty to be added to the model due to the data being linearly separable, which prevented convergence. This means the p-values presented are not technically accurate. For the logistic and Cox models, the test is whether the parameter in question has a non-zero slope, which is impacted by the values of all other features.

Table 1 shows the resulting p-values. The logistic model finds only one significant relationship, which is not corroborated by the data or previous analysis. In addition, it marks several features as unimportant ($p=1.0$) that were previously found to be highly significant. This shows that a classification model of reproducibility is broadly not appropriate for modeling reproducibility. In general, there is considerable agreement between the Cox model and the original non-parametric testing of individual factors. These results in isolation give us an increased confidence that the Cox proportional hazards assumption is reasonable, as it has reasonable correspondence with the original analysis. We note that if one instead

performs the Independent tests on the same sub-set of data (due to aforementioned 72 papers removed), the results are robust. Only the two features *italicized* change, with *Exact Compute* becoming significant and *Code Available* becoming non-significant.

If the proportional hazard assumption is correct, we should observe that the hazard ratio over time for any two points i and j should remain a constant for all time t . This can be tested using the Kaplan–Meier (KM) (Kaplan and Meier 1958) and log-rank (Mantel 1966) tests.

Feature	Transform	Statistic	p-value
Normalized Number of Equations	km	4.30	0.04
	rank	8.03	<0.005
Year Attempted	km	5.78	0.02
	rank	8.76	<0.005

Table 2: Statistical Test of Cox proportional hazard assumption for each feature, shown for the two cases that rejected the null hypothesis.

We tested this assumption for each feature, and found only two that rejected the null hypothesis. This was the Normalized Number of Equations and Year Attempted features, as shown in Table 2. We note that in isolation only two features failing out of 34 coefficients in the model is encouraging, and it corresponds with having $34 \cdot 0.05 = 1.7$ false positives at a significance level of $\alpha = 0.05$. We also investigate these two features further by evaluating their residual fits under the linear model (see appendix). In both cases, we can see that the Cox model is actually a decent fit, with the residual errors concentrated about zero. The failure appears to be due to the linear nature of the current Cox model, which we address in §5.

The agreement between the Cox model on significant factors, as well as the Cox model passing goodness of fit tests, are all indicators of the appropriateness of this approach. As such, it is not surprising that we would obtain reasonable predictive performance with the Cox model. Specifically, we measure the Concordance (Harrell Jr., Lee, and Mark 1996) for the linear model, defined as $\frac{2C+T}{2(C+D+T)}$. Here C refers to the number of pairs where there model correctly predicts the survival order (i.e., which paper was reproduced first is ordered first), D the number of pairs where the model predicted the wrong order, and T the number of ties made by the model. Under Concordance, a random model would receive a score of 0.5. Using 10-fold cross validation, the linear Cox model presented here obtains 0.73.

4.2 Analysis of Linear Cox Model

Given the above results, it is clear that the Cox proportional hazards model is an appropriate method to model the joint set of features and analyze how they impact reproducibility time. The assumptions of the Cox model appear to hold for our data, and the p-values indicate that the model has correspondence with the original analysis as well as superior fit compared to a simpler logistic approach that treats the problem as a binary

Feature	β_i	$\exp(\beta_i)$	std. err
Year	0.00	1.00	0.03
Year Attempted	-0.10	0.91	0.13
Has Appendix	-0.57	0.57	0.31
Uses Exemplar Toy Problem	-0.45	0.64	0.35
Exact Compute Used	0.30	1.35	0.35
Looks Intimidating	0.47	1.61	0.37
Data Available	-0.12	0.89	0.50
Author Code Available	0.42	1.52	0.31
Number of Authors	0.03	1.04	0.09
Pages	0.01	1.01	0.02
Normalized Num References	0.09	1.10	0.15
Normalized Number of Equations	0.02	1.02	0.07
Normalized Number of Proofs	0.42	1.52	0.58
Normalized Number of Tables	-0.13	0.88	0.73
Normalized Number of Graphs/Plots	-0.08	0.92	0.12
Normalized Number of Other Figures	0.02	1.02	1.06
Normalized Conceptualization Figures	1.28	3.59	0.84
Hyperparameters Specified_No	-0.13	0.87	0.25
Hyperparameters Specified_Partial	0.53	1.70	0.76
Hyperparameters Specified_Yes	0.03	1.03	0.20
Paper Readability_Excellent	1.49	4.46	0.37
Paper Readability_Good	0.99	2.70	0.22
Paper Readability_Ok	-0.17	0.85	0.31
Paper Readability_Low	-1.29	0.27	0.24
Algo Difficulty_High	-0.02	0.98	0.23
Algo Difficulty_Medium	-0.42	0.66	0.20
Algo Difficulty_Low	0.44	1.56	0.23
Pseudo Code_Code-Like	-0.03	0.97	0.45
Pseudo Code_Yes	-0.02	0.98	0.20
Pseudo Code_Step-Code	-0.65	0.52	0.35
Pseudo Code_No	0.68	1.97	0.21
Rigor vs Empirical_Balance	0.45	1.57	0.19
Rigor vs Empirical_Empirical	-0.15	0.86	0.35
Rigor vs Empirical_Theory	-0.32	0.73	0.32

Table 3: Coefficients of the linear Cox hazard model.

classification problem.

While we will show later that a non-linear Cox approach better models the data, we find it instructive to first examine the insights and results of the linear approach as it is easier to interpret. In Table 3, we present the coefficients β_i for each variable, its exponentiated value $\exp(\beta_i)$, and the estimated standard error for each coefficient. The value $\exp(\beta_i)$ is of particular relevance, as in the linear Cox model it can be interpreted as the proportional impact on the paper’s reproduction time per unit increase in the feature. For example, this table indicates the Year a paper was published has no impact with a value $\exp(\beta_{\text{Year}}) = 1$.

In Raff (2019), it was found that a paper’s readability was a significant relationship. While the obvious interpretation of these results was that more readable papers are more reproducible, the binary measure of reproducibility used gives us no means to quantify the nature of this relationship. We can now quantify this nature under the framework of survival analysis. Papers that had “Excellent” readability reduced the reproduction time by $4.46\times$ and “Good” papers by $2.70\times$. Those that where “Ok” begin to take longer, increasing the time to reproduce by $0.85^{-1} \approx 1.18\times$ and “Low” readability by $0.27^{-1} \approx 3.70\times$.

In examining the coefficients presented here, we highlight

the importance of considering the standard error. For example, the second most impactful feature appears to be the number of Conceptualization Figures per page, but the standard error is nearly the same magnitude as the coefficient itself. This indicates that the relationship may not be as strong as first indicated by the linear Cox model.

5 Non-Linear Cox and Interpretation

We have now used the linear Cox model to evaluate the appropriateness of a Cox model, where we can more readily compare the Cox Proportional Hazard assumptions with the independent statistical testing of prior work, as well as a logistic alternative. Now, we use a non-linear tree-based model of the Cox assumption and focus our evaluation on the results, and their meaning, with this non-linear model. We do this because the data exhibits non-linear behavior and interactions that made it difficult to fit the linear Cox model. This prevents us from using standard statistical tests on the coefficients with full confidence, and we know that outliers and non-linear behavior plays a non-trivial impact on the results of the Cox regression, $\lambda(t|\mathbf{x}_i) = \exp(f(\mathbf{x}_i)) \cdot \lambda_0(t)$.

In particular, the XGBoost library implements a Cox log-likelihood based splitting criteria (Bou-Hamad, Larocque, and Ben-Ameur 2011; LeBlanc and Crowley 1992), allowing us to use a boosted decision tree to model the $\exp(\mathbf{x}_i^T \beta)$ term of the standard linear Cox model. Optuna (Akiba et al. 2019) was used to tune the parameters of the model resulting in a 10-fold cross-validated Concordance score of 0.80, which is a significant improvement over the linear Cox model indicating a better fit. For this model, we encoded X features as ordinal values that correspond with the nature of what they measure since XGBoost does not support multinomial features. This included the Paper Readability (*Low=0, Ok=1, Good=2, Excellent=3*), Algorithm Difficulty (*Low=0, Medium=1, High=2*), and Pseudo Code (*None=0, Step-Code=1, Yes=2, Code-like=3*) features. All others were encoded in a one-hot fashion.

With a boosted survival tree as our model, we can use SHapley Additive exPlanations (SHAP) (Lundberg and Lee 2017) as a method of understanding our results in greater detail and nuance than the linear Cox model allows. The SHAP value indicates the change in the model’s output contributed by a specific feature value. This is measured for each feature of each datum, giving us a micro-level of interpretability, while we can obtain a macro-level by looking at the average or distribution of all SHAP values. The SHAP value unifies several prior approaches to interpreting the impact of a feature \mathbf{x} ’s values on the output (Ribeiro, Singh, and Guestrin 2016; Štrumbelj and Kononenko 2014; Shrikumar, Greenside, and Kundaje 2017; Datta, Sen, and Zick 2016; Bach et al. 2015; Lipovetsky and Conklin 2001). In addition, Lundberg, Erion, and Lee (2018) showed how to exactly compute a tree’s SHAP score in $O(LD^2)$ time, where L is the maximum number of leaf nodes and D the maximum depth. The Tree SHAP is of particular value because it allows us to disentangle the impact of feature interactions within the tree, giving us a feature importance that makes it easier to understand on an individual basis with less complication.

A simple plot of all the SHAP values, sorted by their importance, is presented in Figure 3. Red indicates higher per-feature values and blue lower. The x-axis SHAP values are changes in the log hazard ratio. Positive SHAP values correspond to features decreasing the reproduction time, while negative values indicate the feature correlates with an increase the time to reproduce. For example, a SHAP score of 0.4 would be a $\exp(0.4) \approx 1.49 = 49\%$ reduction in the time to reproduce. At a high level, we can immediately see that a number of the features identified as important in Raff (2019) are also important according the Cox model, with the same overall behaviors. Paper Readability, Number of Equations, Rigor vs Empirical, Algorithm Difficulty, Hyperparameters Specification, and Pseudo-Code are all important. Other features such as the Use of an Exemplar Toy Problem, Having an Appendix, Data Made Available, or Looking Intimidating have no significant impact.

We also see what may at first appear to be contradictions with prior analysis. The number of Graphs/Plots was originally determined to have a non-significant relationship. In our model having fewer Graphs/Plots appears to be one of the most significant factors. We explain this by again noting the distinction of our model: the Cox modeling presented above indicates the impact on *how long* it will take us to reproduce the result, whereas the result by Raff (2019) indicates only if the feature has a significant impact on the *success* of reproduction. As such, we argue there is no contradiction in the results but only deeper insight into their nature. In the case of Graph/Plot, we say that the number of graphs or plots does not impact whether a paper can be reproduced; however, if it is reproducible then having too many graphs and plots increases the time it takes to reproduce.

Due to space limitations, we can not discuss all of the behaviors that are observable from this data. For this reason, we will attempt to avoid discussing results that reaffirm previous conclusions. We will discuss features in this context if the new Cox analysis provides a particular new insight or nuance into the nature of the relationship between said feature and reproducibility.

We note that most categorical features have a relatively simple analysis, and the results previously found to be significant also have the most impact on reproduction time. For example, one can clearly see that *Excellent* paper readability decreases reproduction time while *Low* readability increases it. Results are also consistent with the sign of correlation previously found, e.g., *Theory* papers take more time to reproduce than *Empirical* ones. However, new insight is derived from noting that *Balanced* papers further reduce the reproduction time.

Of these categorical variables, we draw special attention to three results. First, the nature of including pseudo-code is surprising. Before, it was noted that *Step-Code* negatively correlated with reproducibility, but papers with *No* and *Code-Like* pseudo-code were more reproducible. Under the Cox model, papers with *No* pseudo-code take significantly less time to reproduce. *Step-Code* is found to have a negative impact, while simply having pseudo-code or *Code-Like* pseudo-code have almost no impact on the hazard ratio.

We also draw attention to the fact that specifying the Exact Compute Used and the authors making their code available

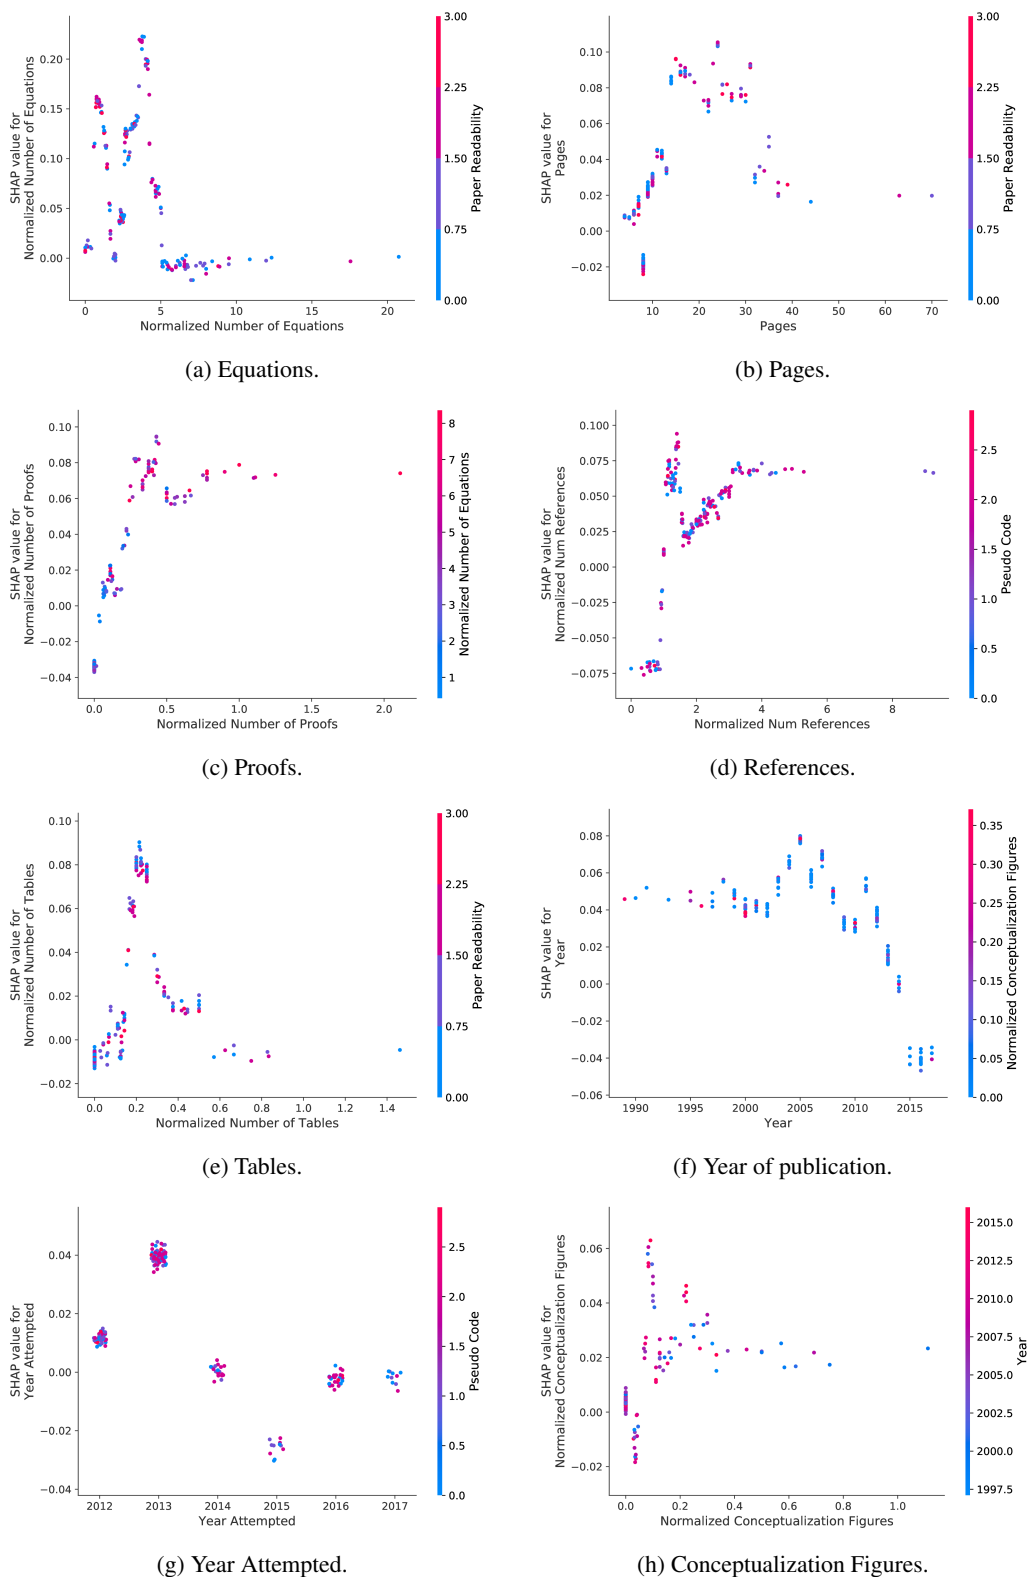


Figure 2: SHAP individual features (change in log-hazard ratio) for several numeric features on the y-axis with feature values on the x-axis. Each figure has a color set based on the value of a second feature, indicated on the right. The second feature is selected by having the highest SHAP interaction.




Figure 3: SHAP results for each feature. The change in log hazard ratio (x-axis) caused by each feature where the color coding indicates the values from high (red) to low (blue).

lead to reductions in compute time. As part of the original study, all papers with Code Available were excluded if the code was used prior to reproduction. We suspect that the reduction in reproduction time is a psychological one: that more effort is expended with higher confidence that the results can be replicated when the authors make their code available. A similar hypothesis might be drawn by specifying the Exact Compute Used leading to a reduced reproduction time. Alternatively, it may be the case that the exact equipment provides useful information for the reproducer to judge the meaning behind discrepancies in performance and to adjust accordingly. For example, if the nature of a proposed speedup is dependent on reducing communication overhead between CPU threads, and the reproducer has a system with slower per-core clock speeds than the paper, it could explain getting reduced results — and that final replication may need to be tested on different equipment.

Detailed Analysis of Numerical Features: It was originally noted that the features determined to be significantly related to reproducibility were also the most subjective, which reduces the utility of the results. Under the boosted Cox model, we can more easily and readily derive insights from the more objective features. In Figure 2, we provide individ-

ual plots of all the numeric features. In each case, the SHAP value is plotted on the y-axis and a color is set based on the value of a second feature. The second feature is selected based on the SHAP interaction scores, where the feature with the highest mean interaction score is used to color the plot. This provides additional insights for some cases. However, the degree of non-linear interaction between features is relatively small so the coloring is not always useful.

With respect to the length of a paper (Figure 2b), we see that there is a strong positive correlation with paper length, with most Journal papers having reduced reproduction time. This relationship ends with excessively long papers (≥ 35 pages) that return to the base hazard rate. We believe this relationship stems from the observation that page limits may discourage more implementation details (Raff 2019). Under the Cox model, it becomes clearer that larger page limits in conferences, or more emphasis on Journal publications, is of value in reducing the time it takes to reproduce research.

We anticipate the nature of having longer papers reducing reproduction times is similar to that of the number of references included in the paper (Figure 2d), where a strong positive relationship exists but instead plateaus after ≥ 4 references per page. In the case of References, having too few introduces an increase in reproduction time. Beyond providing empirical evidence for good scholastic behavior, references also serve as valuable sources of information for the techniques being built upon, allowing the reproducer to read additional papers that may further benefit their understanding of the paper they are attempting to reproduce.

It was previously noted that Tables had a significant relationship with the ability to reproduce a paper while Graphs did not. The Cox model does not contradict this conclusion from Raff (2019), but it does refine our understanding. Having no, or less than one graph per page, corresponds to a lower reproduction time, and ≥ 1 Graphs per page results in a minor increase in time. This appears to indicate that Graphs, while visually pleasing, may not be as effective of a visualization tool as the author might hope. We suspect that further delineating the nature and type of graphs would be valuable before drawing broader conclusions. Tables, by comparison, predominantly help reproducibility. However, there is a narrow range in the number of tables per page that leads to a reduction in reproduction time, after which the value decays to become negligible.

Of particular interest are the number of proofs (Figure 2c) and equations (Figure 2a) due to their natural relationship with *Theory* oriented papers. A new insight is that the more proofs in a paper the less time it takes to reproduce its results, but this value reaches a plateau near one proof for every two pages. This gives additional credence to arguments proposed by Sculley et al. (2018) on the importance of understanding the methods we develop at a scientific level. As can be seen in Figure 2c, more proofs does naturally lead to more equations per page, which Figure 2a show has a non-linear relationship. Having no equations is bad, but having up to five per page is connected with a larger reduction in reproduction time than the benefit of more complete proofs. We note, however, that based on the protocol used it may be possible to obtain the benefits of both a moderate number of equations and

well-proven works. The number of equations is counted only within the content of the main paper, but proofs are still counted when they appeared in appendices (Raff 2019). This suggests a potential ideal strategy for organizing the content of papers: few and judicious use of equations within the main paper to provide the necessary details for implementation and replication with more complete examination of the proofs in a separate section of the appendix. This we note should not be extrapolated to works that are purely theory based, which were beyond the scope of the original data.

Last, we note some unexpected results under the Cox model compared to the original analysis. First, in Figure 2f we can see that the amount of time to reproduce a paper has been increasing since 2005. This gives credence to the concern that we are entering a reproducibility crisis within machine learning, as argued by Hutson (2018) which Raff (2019) did not find evidence of.

6 Conclusion

By extending a recent dataset with “reproduction times,” measuring the number of days it took to reproduce a paper, we have proposed and evaluated the first survival analysis of machine learning research reproducibility. Combining this with more modern tree ensembles, we can use SHAP interpretability scores to gain new insights about what impacts the time it will take to reproduce a paper. In doing so, we obtain new insights into the utility of the most objective features of papers while also deriving new nuance to the nature of previously determined factors. We hope this will encourage others in this space to begin recording time spent on reproduction efforts as valuable information.

Acknowledgements

I’d like to thank Drew Farris, Frank Ferraro, Cynthia Matuszek, and Ashley Klein for their feedback and advice on drafts of this work. I would also like to thank Reviewer #3 of my NeurIPS paper for spawning this follow-up work. I woefully underestimated how long it would take and the analyze “how long it took to reproduce”. I hope this paper is a satisfying followup.

Broader Impact

Our work as the potential to positively influence the paper writing and publishing process, and in the near term, better equip us to study how reproducibility can be quantified. The timing of effort spent is intrinsically more objective than “success/non-success,” as it does not preclude any failed attempt from eventually becoming successful. This is important at a foundational level to ensure machine learning operates with a scientific understanding of our results, and to understand the factors that may prevent others from successful replication. Indeed, by making the measure of reproducibility a more objective measure of time spent (though itself not perfect), we may increase the reproducibility of research studies.

With these results come an important caveat that has been emphasized in the paper; that these results do not have exact timing information but rather a proxy measure, and we must

use the survival model to circumvent missing information from the failed replication cases. Our analysis also can not explore all possible alternative hypothesis of the results. We risk readers taking these statements as a dogmatic truth and altering their research practices for the worse if it is not fully understood that this work provides the foundation for a new means of studying reproducibility.

Our analysis must not be taken as a final statement on the nature of these relationships but instead as initial evidence. There are many confounding and unobserved factors that will impact the results and cannot be obtained. It is necessary that the reader understand these to be directions for future research with thought and care to design study protocols that can unravel these confounders. Further still, this study’s data represents only attempts to obtain the same results purported in prior works. This is distinct from whether the conclusions from prior papers may be in error due to insufficient effort at tuning baselines or other factors. This is a distinct and important problem beyond the scope of this work, but has been found relevant to a number of sub-domains (Dacrema, Cremonesi, and Jannach 2019; Blalock et al. 2020; Musgrave, Belongie, and Lim 2020).

The caveats are amplified because the study is based on work with one reproducer. As such, it may not generalize to other individuals, and it may not generalize to other sub-disciplines within machine learning. To do so would require significant reproduction attempts from many individuals with varied backgrounds (e.g. training, years experience, education) and with careful controls. The original study was over a span of 8 years, so progress is unlikely unless a communal effort to collect this information occurs. A cohort of *only 125 participants* would constitute a *millennia of person effort* to make a better corpus over the same conditions. Our hope is that this work, by showing how we can better quantify reproduction attempts via time spent, gives the requisite objective needed to begin such larger communal efforts.

References

- Akiba, T.; Sano, S.; Yanase, T.; Ohta, T.; and Koyama, M. 2019. Optuna: A Next-generation Hyperparameter Optimization Framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD ’19, 2623–2631. New York, NY, USA: ACM. ISBN 978-1-4503-6201-6. doi:10.1145/3292500.3330701.
- Bach, S.; Binder, A.; Montavon, G.; Klauschen, F.; Müller, K.-R.; and Samek, W. 2015. On Pixel-Wise Explanations for Non-Linear Classifier Decisions by Layer-Wise Relevance Propagation. *PLOS ONE* 10(7). ISSN 1932-6203. doi: 10.1371/journal.pone.0130140.
- Barba, L. A. 2019. Praxis of Reproducible Computational Science. *Computing in Science & Engineering* 21(1): 73–78. ISSN 1521-9615. doi:10.1109/MCSE.2018.2881905.
- Blalock, D.; Gonzalez Ortiz, J. J.; Frankle, J.; and Gutttag, J. 2020. What is the State of Neural Network Pruning? In *Proceedings of Machine Learning and Systems 2020*, 129–146.
- Bou-Hamad, I.; Larocque, D.; and Ben-Ameur, H. 2011. A

- review of survival trees. *Statist. Surv.* 5: 44–71. ISSN 1935-7516. doi:10.1214/09-SS047.
- Bouthillier, X.; Laurent, C.; and Vincent, P. 2019. Unreproducible Research is Reproducible. In Chaudhuri, K.; and Salakhutdinov, R., eds., *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, 725–734. Long Beach, California, USA: PMLR, URL <http://proceedings.mlr.press/v97/bouthillier19a.html>.
- Collaboration, O. S. 2015. Estimating the reproducibility of psychological science. *Science* 349(6251). ISSN 0036-8075. doi:10.1126/science.aac4716.
- Dacrema, M. F.; Cremonesi, P.; and Jannach, D. 2019. Are we really making much progress? A Worrying Analysis of Recent Neural Recommendation Approaches. In *Proceedings of the 13th ACM Conference on Recommender Systems - RecSys '19*, 101–109. New York, New York, USA: ACM Press. ISBN 9781450362436. doi:10.1145/3298689.3347058.
- Datta, A.; Sen, S.; and Zick, Y. 2016. Algorithmic Transparency via Quantitative Input Influence: Theory and Experiments with Learning Systems. In *2016 IEEE Symposium on Security and Privacy (SP)*, 598–617. IEEE. ISBN 978-1-5090-0824-7. doi:10.1109/SP.2016.42.
- Forde, J.; Head, T.; Holdgraf, C.; Panda, Y.; Perez, F.; Nalvarte, G.; Ragan-kelley, B.; and Sundell, E. 2018. Reproducible Research Environments with repo2docker. In *Reproducibility in ML Workshop, ICML'18*.
- Geburu, T.; Morgenstern, J.; Vecchione, B.; Vaughan, J. W.; Wallach, H.; Daumeé, H.; and Crawford, K. 2018. Datasheets for Datasets. *ArXiv e-prints* 1–27, URL <http://arxiv.org/abs/1803.09010>.
- Glenn, B. C.; and P.A., I. J. 2015. Reproducibility in Science: Improving the Standard for Basic and Preclinical Research. *Circulation Research* 116(1): 116–126. doi:10.1161/CIRCRESAHA.114.303819.
- Gronenschild, E. H. B. M.; Habets, P.; Jacobs, H. I. L.; Mengelers, R.; Rozendaal, N.; van Os, J.; and Marcelis, M. 2012. The Effects of FreeSurfer Version, Workstation Type, and Macintosh Operating System Version on Anatomical Volume and Cortical Thickness Measurements. *PLOS ONE* 7(6): 1–13. doi:10.1371/journal.pone.0038234.
- Gundersen, O. E.; and Kjensmo, S. 2018. State of the Art: Reproducibility in Artificial Intelligence. *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI-18)* 1644–1651.
- Harrell Jr., F. E.; Lee, K. L.; and Mark, D. B. 1996. Multivariable Prognostic Models: Issues in Developing Models, Evaluating Assumptions and Adequacy, and Measuring and Reducing Errors. *Statistics in Medicine* 15(4): 361–387. ISSN 0277-6715. doi:10.1002/(SICI)1097-0258(19960229)15:4<361::AID-SIM168>3.0.CO;2-4.
- Hutson, M. 2018. Artificial intelligence faces reproducibility crisis. *Science* 359(6377): 725–726. ISSN 0036-8075. doi:10.1126/science.359.6377.725.
- Ioannidis, J. P. 2017. The Reproducibility Wars: Successful, Unsuccessful, Uninterpretable, Exact, Conceptual, Triangulated, Contested Replication. *Clinical Chemistry* 63(5): 943–945. ISSN 0009-9147. doi:10.1373/clinchem.2017.271965.
- Ioannidis, J. P. A. 2018. Meta-research: Why research on research matters. *PLOS Biology* 16(3): e2005468, URL <https://doi.org/10.1371/journal.pbio.2005468>.
- Kaplan, E. L.; and Meier, P. 1958. Nonparametric Estimation from Incomplete Observations. *Journal of the American Statistical Association* 53(282): 457–481. ISSN 0162-1459. doi:10.1080/01621459.1958.10501452.
- LeBlanc, M.; and Crowley, J. 1992. Relative Risk Trees for Censored Survival Data. *Biometrics* 48(2): 411. ISSN 0006341X. doi:10.2307/2532300.
- Lin, D. Y.; and Wei, L. J. 1989. The Robust Inference for the Cox Proportional Hazards Model. *Journal of the American Statistical Association* 84(408): 1074–1078. ISSN 01621459. doi:10.2307/2290085.
- Lipovetsky, S.; and Conklin, M. 2001. Analysis of regression in game theory approach. *Applied Stochastic Models in Business and Industry* 17(4): 319–330. ISSN 1524-1904. doi:10.1002/asmb.446.
- Lundberg, S. M.; Erion, G. G.; and Lee, S.-I. 2018. Consistent Individualized Feature Attribution for Tree Ensembles. *ArXiv e-prints* (2), URL <http://arxiv.org/abs/1802.03888>.
- Lundberg, S. M.; and Lee, S.-I. 2017. A Unified Approach to Interpreting Model Predictions. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 30*, 4765–4774. Curran Associates, Inc. URL <http://papers.nips.cc/paper/7062-a-unified-approach-to-interpreting-model-predictions.pdf>.
- Mantel, N. 1966. Evaluation of survival data and two new rank order statistics arising in its consideration. *Cancer chemotherapy reports* 50(3): 163–170. ISSN 0069-0112, URL <http://europepmc.org/abstract/MED/5910392>.
- Mesnard, O.; and Barba, L. A. 2017. Reproducible and Replicable Computational Fluid Dynamics: It's Harder Than You Think. *Computing in Science & Engineering* 19(4): 44–55. ISSN 1521-9615. doi:10.1109/MCSE.2017.3151254.
- Musgrave, K.; Belongie, S.; and Lim, S.-N. 2020. A Metric Learning Reality Check. *arXiv* URL <http://arxiv.org/abs/2003.08505>.
- Olorisade, B. K.; Brereton, P.; and Andras, P. 2018. Reproducibility in Machine Learning-Based Studies : An Example of Text Mining. In *Reproducibility in ML Workshop, ICML'18*.
- Prinz, F.; Schlange, T.; and Asadullah, K. 2011. Believe it or not: how much can we rely on published data on potential drug targets? *Nature Reviews Drug Discovery* 10(9): 712. ISSN 1474-1784. doi:10.1038/nrd3439-c1.
- Raff, E. 2017. JSAT: Java Statistical Analysis Tool, a Library for Machine Learning. *Journal of Machine Learning Research* 18(23): 1–5, URL <http://jmlr.org/papers/v18/16-131.html>.

Raff, E. 2019. A Step Toward Quantifying Independently Reproducible Machine Learning Research. In *NeurIPS*. URL <http://arxiv.org/abs/1909.06674>.

Ribeiro, M. T.; Singh, S.; and Guestrin, C. 2016. “Why Should I Trust You?”: Explaining the Predictions of Any Classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16, 1135–1144. New York, NY, USA: ACM. ISBN 978-1-4503-4232-2. doi:10.1145/2939672.2939778.

Sculley, D.; Holt, G.; Golovin, D.; Davydov, E.; Phillips, T.; Ebner, D.; Chaudhary, V.; Young, M.; Crespo, J.-F.; and Densnison, D. 2015. Hidden Technical Debt in Machine Learning Systems. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2*, NIPS'15, 2503–2511. Cambridge, MA, USA: MIT Press.

Sculley, D.; Snoek, J.; Rahimi, A.; and Wiltschko, A. 2018. Winner’s Curse? On Pace, Progress, and Empirical Rigor. In *ICLR Workshop track*. ISBN 1559-713X, URL <https://openreview.net/pdf?id=rJWF0Fywf>.

Shrikumar, A.; Greenside, P.; and Kundaje, A. 2017. Learning important features through propagating activation differences. In *34th International Conference on Machine Learning, ICML 2017*, 4844–4866. ISBN 9781510855144.

Štrumbelj, E.; and Kononenko, I. 2014. Explaining Prediction Models and Individual Predictions with Feature Contributions. *Knowl. Inf. Syst.* 41(3): 647–665. ISSN 0219-1377. doi:10.1007/s10115-013-0679-x.

Wicherts, J. M.; Borsboom, D.; Kats, J.; and Molenaar, D. 2006. The poor availability of psychological research data for reanalysis. *American Psychologist* 61(7): 726–728. ISSN 1935-990X. doi:10.1037/0003-066X.61.7.726.