

Written in the Stars: How your (pens and) papers decide the fate of the arXiverse

JOANNE TAN^{1,*}

¹*Earth, Maybe. Perhaps an Extragalactic Existence* , or the Equivalent of a Hocus-pocus (EMPEE-EH).[†]

ABSTRACT

We all *love* the ecstasy that comes with submitting papers to journals or arXiv. Some have described it as yeeting their back-breaking products of labor into the void, wishing they could never deal with them ever again. The very act of yeeting papers onto arXiv contributes to the expansion of the arXiverse; however, we have yet to quantify our contribution to the cause. In this work, I investigate the expansion of the arXiverse using the arXiv **astro-ph** submission data from 1992 to date. I coin the term “the arXiverse constant”, a_0 , to quantify the rate of expansion of the arXiverse. I find that **astro-ph** as a whole has a positive a_0 , but this does not always hold true for the six subcategories of **astro-ph**. I then investigate the temporal changes in a_0 for the **astro-ph** subcategories and **astro-ph** as a whole, from which I infer the fate of the arXiverse. I conclude that our arXiverse is past its peak of expansion and could gradually slow down to a crunch.

1. INTRODUCTION

arXiv is an open-access e-print archive that hosts multi-million scholarly articles in various fields. Researchers regularly submit their articles to arXiv as a way to share their research with people around the world. arXiv host articles from several main fields, namely physics, mathematics, computer science, quantitative biology, quantitative finance, statistics, electrical engineering and systems science, and economics. Each of these scientific fields also consists of multiple sub-fields. arXiv was founded in August 1991 and has been growing steadily ever since. Honestly, if you are reading this paper, I don’t think I need to explain arXiv in more detail — you are already on arXiv! 😊

While arXiv has been growing in popularity among researchers, there has not been an effort to quantify their contribution to the growth (or expansion) of arXiv. I find this intriguing since we researchers **LOVE** to submit our papers into the void of arXiv (and of journals, of course), but we never really stop to think about the consequences of our seemingly mundane and harmless action. We have derived so much euphoria from yeeting our paper to arXiv, it is only right that we also consider the effects of our action. Therefore, in this paper, I aim to quantify our contribution to the expansion of

the arXiv. With that aim in mind, I set off to hunt down arXiv’s submission statistics, but only from the **astro-ph** category. This choice is motivated by two reasons: (i) it is the main category I yeet my research to; (ii) it was set up shortly after the birth of arXiv and hence the data spans a temporal baseline of $\simeq 32$ years. In §2, I describe the methods I have utilized to obtain the dataset necessary for this work, and the linear regression fits I perform on the datasets. In §3, I discuss the rate of expansion of the arXiv inferred from the model fits and the fact that it is not constant in time. Then, in §4, I show robust proof that the rate of expansion changes over time, and discuss the limitations of this work, before delving into the future prospects of this work. I conclude the findings in §5.

2. METHODS

2.1. *Yeinking the required data*

I¹ utilize two query methods to obtain the data I need for this work. Firstly, I query the arXiv via its API² for papers submitted to arXiv **astro-ph** from 1992 to 2008. I heavily modified the Python package **arxivscraper**³(Sadjadi 2017) to query for submitted papers using the arXiv API instead of the arXiv OAI-PMH interface⁴. I filter for submissions to the **astro-ph**

* The author is responsible for the contents of this paper, which do not in any way represent the views of their employers. Their views are their own, no matter if scrutinized through a microscope or a telescope.

[†] Read the acronym aloud to guesstimate where the author is based.

¹ Using the singular I in this paper because using we sounded really off in my head while writing.

² <https://info.arxiv.org/help/api/index.html>

³ <https://github.com/Mahdisadjadi/arxivscraper>

⁴ <https://info.arxiv.org/help/oa/index.html>

category from the queried metadata. Then, I select only the submissions with `astro-ph` as their primary category — papers submitted to `astro-ph` as a secondary category are all excluded from the dataset. For every month from January 1992 to December 2008, I compute the number of papers submitted to `astro-ph` as their primary category and the number of unique submission dates. With these, I calculate the average daily submission rate for any given month.

The second query method was using `arxivscraper` directly to query arXiv for submission data from 2009 onwards. This was mainly because the first method was not able to handle the queries adequately, resulting in a time-consuming process to query the required metadata. The temporal split of the dataset in 2009 is due to the fact that `astro-ph` was growing exponentially and arXiv decided to introduce six subcategories for `astro-ph` namely `astro-ph.GA`, `astro-ph.SR`, `astro-ph.CO`, `astro-ph.EP`, `astro-ph.HE`, and `astro-ph.IM`, in January 2009. `arxivscraper` adequately handles queries efficiently for data since 2009+, for separate subcategories. Here, I apply a similar series of steps as the first method. I filter for submissions to the selected `astro-ph` subcategory, and then select only those that have the selected subcategory as their primary category. Then, for every month from January 2009 to February 2024, I calculate the number of papers submitted to the particular `astro-ph` subcategory and the number of unique submission dates. I then compute the average daily submission rate for the subcategory. To get the submission statistics for `astro-ph` as a whole, I sum up the number of papers submitted across the six subcategories to get the total submitted papers for the month. On the other hand, I use the number of days in the selected month as the number of unique submission dates. This decision is motivated by the fact that it is statistically very likely that at least one paper is submitted to `astro-ph` on any given day.⁵ To ensure the robustness of this method, I compare the mined data to arXiv’s very own submissions statistics⁶ (available for submissions from 2009 onwards) and also a mini query using the first method for `astro-ph.GA`. The comparisons show the data from the three sources/queries are highly similar. Therefore, despite the two different methods to query for submission statistics, I deduce that the two datasets can be combined temporally. From the queries, I obtain submission statistics for the following eight data groups:

1. `astro-ph.GA` since 2009+
2. `astro-ph.SR` since 2009+
3. `astro-ph.CO` since 2009+
4. `astro-ph.EP` since 2009+
5. `astro-ph.HE` since 2009+
6. `astro-ph.IM` since 2009+
7. `astro-ph` since 2009+ (combination of groups 1–6)
8. `astro-ph` since 1992+ (data from first query method, combined with group 7)

2.2. Fitting models to the data

I perform a linear regression (LR) fitting onto the average daily submission rate as a function of time, by using `sklearn.linear_model.LinearRegression` and `scipy.stats.linregress`. I first fit the LR model onto the complete dataset of a group (as defined at the end of Section 2.1) to obtain the slope and intercept of the LR line. Hereafter, I refer to this fit as LR-full. Then, I split the dataset into a training dataset and a testing dataset using the 80:20 split. I fit an LR model onto the training dataset and then use the model to generate predictions using the testing dataset. This fit is referred to as LR-split hereafter. Next, I compare the slopes and intercepts of the LR fits on the full dataset and split dataset obtained from both `sklearn.linear_model.LinearRegression` and `scipy.stats.linregress`; the results are found to be identical up to at least 6 s.f. Hence, I show only the results from `scipy.stats.linregress` as this function also provides the standard errors of the slopes and intercepts. Last but not least, I determine the 95% confidence interval of the LR fits. At every time point, I draw 1000 random samples from a Gaussian distribution, with the slope as the center and the slope’s standard error as the width. I repeat the same random draws for the intercept and its standard error. Every draw provides a predicted submission rate, giving me a distribution of 1000 predictions at every point in time. I then compute the 2.5th and 97.5th percentile of the distribution to form the 95% confidence interval of the fit.

Figure 1 shows the LR-full and LR-split fits for the submission data of `astro-ph.EP` (top panel) and of `astro-ph.HE` (bottom panel). Upon inspection, one notices that the slope of both LR-full fits is slightly smaller than that of the LR-split fits. I infer that the submission rate has increased slower in recent times. This is due to the fact that the LR-split is essentially a fit onto 4/5 of the dataset but applies the model onto the full dataset. Since the slope decreases when the LR model considers data from recent times, I deduce that the rates of expansion of these two `astro-ph` subcategories are decreasing.

⁵ Yes, this is very handwavy, but ~20-31 days are on the same order of magnitude anyway.

⁶ https://info.arxiv.org/about/reports/submission_category_by_year.html




Figure 1. Submission rate as a function of time for `astro-ph.EP` (group 4, top panel) and `astro-ph.HE` (group 5, bottom panel). In both panels, the purple stars indicate the actual submission rate in a given month. The orange line indicates the LR fit from the full dataset (LR-full), while the cyan dashed line indicates the LR fit from the split dataset (LR-split). The shaded regions around the LR fits indicate the 95% confidence interval of the LR fit. The difference in the slopes of the LR-full and LR-split fits indicates a change in the expansion rate of the arXiverse.

Unfortunately, the confidence intervals do not include most of the actual submission data. I will discuss this observation in Section 4 and the possible explanations. Let’s now shift our focus to Figure 2, which shows both the LR-full and LR-split fits for the aggregate submission data for `astro-ph` since 2009+ (top panel) and for `astro-ph` since 1992+.

deduction is made — the rates of expansion of these two data groups are decreasing over time. Additionally, as I am using `astro-ph` as a whole to be a proxy of the arXiverse, this leads to the inference that the rate of expansion of the arXiverse is decreasing.




Figure 2. Same as Figure 1, but for `astro-ph` in aggregate since 2009+ (group 7, top panel) and for `astro-ph` as a whole since 1992+ (group 8, bottom panel).

3. THE arXiverse CONSTANT

The time is nigh to introduce the backbone of this work — the arXiverse constant, a_0 . Essentially, it is defined as the rate of expansion of the arXiverse. The slope from the LR fits delineated in Section 2.2 describes the (monthly) rate of change in submission rate, and serves

as the (almost⁷) perfect proxy for the rate of expansion of the arXiverse, with units of paper/day/month. The slope from the LR fits for the **astro-ph** subcategories describes the rate of expansion of the respective sub-arXiverse, while the slope from the LR fits for **astro-ph** as a whole would then be the proxy for the arXiverse.

Thus, for each data group, I have two a_0 measurements from the two LR fits (refer to Section 2.2). As I want to have another robust measurement of a_0 for each data group, I implement the Theil-Sen estimator (using `scipy.stats.theilslopes`) to compute the slope, the intercept, and the confidence interval of the slope of a robust LR fit to the complete dataset for a selected group. This method is referred to as TS-full hereafter. Figure 3 compares the a_0 measured from the three LR fits. We can see that the a_0 measured from LR-full and TS-full agree very well for all eight groups. In contrast, LR-split only agrees with LR-full (or TS-full) for four **astro-ph** subcategories and also **astro-ph** in aggregate from 2009 onwards. The lack of consistent agreement is likely due to the split in datasets into training and testing sets. This also indicates that the measured a_0 changes over time. As observed in Figure 2 where the slope measured and thus a_0 decreases over time, this is confirmed in Figure 3, where the a_0 for LR-full for both **astro-ph** in aggregate since 2009+ and for **astro-ph** in aggregate since birth (i.e., 1992+) is smaller than that for LR-split. While it may be very appealing to immediately jump to the conclusion that arXiverse is then experiencing a deceleration in its rate of expansion, I thought it prudent to proceed with caution. Upon a quick glance, we can also observe that except for **astro-ph.CO** and **astro-ph.SR**, all data groups have positive a_0 measured. This indicates that while other **astro-ph** subcategories and even **astro-ph** as a whole may experience an increase in submission rate across time, this is not the case for **astro-ph.CO** and **astro-ph.SR**. These subcategories instead experience a decrease in submission rate and an almost constant submission rate, respectively. Therefore, we can deduce that not every sub-arXiverse experiences the same rate of expansion, let alone the same change in rate of expansion. Thus, I conclude that there is a tension in the arXiverse constant, which can vary depending on the dataset used for measurement.

At this stage, I think it is imperative to draw a clear line between the arXiverse constant and a similar-sounding constant that we are probably very familiar

with, i.e., the Hubble constant, H_0 . While it may seem to readers that I am trying to draw a parallel comparison between the arXiverse and the Universe, I am, in fact, not doing that. The arXiverse constant is *not* meant to be analogous to the Hubble constant as they are measures of distinct “-verses” and are thus not comparable. Any coincidence is purely unintended⁸. Anyway, here is a cartoon illustration for you 😊 In Figure 4, I show how we should feel about the different tensions in cosmic measurements.




Figure 3. The arXiverse constant, a_0 , as measured using three different linear regression fittings on all data groups (labelled on y-axis). The teal diamonds indicate a_0 measured using the `scipy.stats.linregress` method on the full dataset, the purple pentagons indicate a_0 measured using also `scipy.stats.linregress` but on the split dataset, and the orange stars indicate a_0 measured using the Theil-Sen estimator on the full dataset.

4. SCRUTINIZING THE RESULTS

4.1. How does a_0 really change over time? 🤔

In the preceding section, I establish that the a_0 changes over time and there is thus a tension in a_0 measurements. One way to verify this is to calculate the value of a_0 across time, since the birth of the arXiverse (and the sub-arXiverses). For a given data group, I start with 12 data points, i.e., the first 12 months of submission data (average submitted paper per day), and implement the Theil-Sen estimator onto the selected data points to compute the slope and its confidence interval. The calculated slope and confidence interval then serve as the measured a_0 and its confidence interval at that point in time, which in this case, is the 13th month since I calculate the a_0 using all the available data up to that point in time. I then repeat the same process in the subsequent months, adding one data point to the Theil-Sen estimator at every iteration. This will then show us ex-

⁷ Including this in case someone comes at me and says that nothing is perfect, so here’s to me covering all bases.

⁸ Perhaps, maybe. :>



Figure 4. How we should feel about the arXiverse tension. 🤪

actly how a_0 changes with time. I start with a minimum of 12 data points instead of fitting with fewer data points because I think having a year’s worth of data is crucial for a robust initial measurement of a_0 . For data groups 1–7, since the submission rate data starts from January 2009, this means that the first a_0 estimation was done in January 2010. For data group 8, the first estimation was done in April 1993. The a_0 measured across time using this method is shown in the top panel of Figure 5. I have split the `matplotlib` colormap ‘Spectral’ into numerous hues to color-code the a_0 measurements across time. The colormap runs from red to blue hues akin to a rainbow’s spectrum — the red end of the colormap indicates earlier times (ca. 1992), while the blue end indicates later or recent times (ca. 2024). The midpoint of the spectrum is colored by yellow hues, which indicate half of arXiv’s lifetime to date. Coincidentally, the midpoint is around 2009, which was the point in time when `astro-ph` split into 6 subcategories and where we started having submission rate statistics for data groups 1–7.

The top panel of Figure 5 shows clearly that a_0 was rarely constant in time. Now, let’s look into the values of a_0 over time for individual data groups. For `astro-ph.GA` and `astro-ph.SR`, there is an initial monotonic increase in a_0 measurements over time (stars with green hues having larger a_0 values compared to stars with yellow hues) and then a monotonic decrement in a_0 (indicated by stars with blue hues having smaller a_0 values than stars with green hues). In contrast, `astro-ph.CO`

shows a monotonic initial decrease, as indicated by the decrease in a_0 from the yellow-hued stars, to light-green stars, then to green stars. In more recent times, a_0 measured shows a monotonic increment, as indicated by the increase in a_0 from the green stars to the blue stars. These three subcategories solidify the notion that a_0 is not a constant in time. The three remaining subcategories, `astro-ph.EP`, `astro-ph.HE`, and `astro-ph.IM`, as well as `astro-ph` aggregate since 2009+, initially have large fluctuations around the latest measurement of a_0 , which gradually decreases over time, to small fluctuations around the final measurement. They show that while a_0 is not constant over time, it seems to settle in around a “Goldilocks Zone” towards recent times. Lastly, when `astro-ph` is studied as a whole since its birth, it shows similar trends in a_0 as `astro-ph.GA` and `astro-ph.SR` – there is an initial monotonic increment up to a certain point (stars with orange hues have larger values compared to stars with red hues) before a_0 starts to decrease monotonically (a_0 decreases from stars with orange hues to stars with blue hues).

I then simplify the a_0 measurement over time by measuring every subsequent 12 months instead of every subsequent month. This way, the potential seasonal effect in a year, which may cause fluctuations in a_0 measurements to a certain extent, can be absorbed into the LR fits. The a_0 measurements are shown in the bottom panel of Figure 5. The color codings are the same as the top panel. The trends in a_0 observed in the top panel for the eight data groups are also observed to be similar in the bottom panel, except that they are more evident in the bottom panel.

I then look into another metric of a_0 that could be useful, i.e., the annual change in a_0 , or Δa_0 . This is essentially the year-to-year difference in the measured a_0 using the available data up to the selected point in time. The result is shown in Figure 6. The top panel of Figure 6 shows Δa_0 for the full temporal baseline – from 1992 to date (from 2009 to date for data groups 1–7). I observe that all data groups, at the earlier times of their respective dataset, have relatively large fluctuations in a_0 compared to those of recent times. The bottom panel of Figure 6 shows the zoomed-in version of the top panel, focusing on 2009 onwards at a smaller y-axis range. The hovering of Δa_0 around zero from 2017 onwards for all the data groups (and from ca. 2000 onwards for group 8, as seen in the top panel) shows that a_0 is very likely to settle around a “Goldilocks” value soon. It also suggests that about 8 years’ worth of data is required for a more robust measurement of a_0 , as that is the point of time where Δa_0 starts being relatively smaller and tapers off. Upon closer inspection, the




Figure 5. The arXiv constant, a_0 , measured at a given point in time (from January 2010 for data groups 1–7, from April 1993 for group 8, see text for more information regarding dates, groups as defined in Section 2.1), using submission rate data up to selected point in time. Top panel shows a_0 measured every month while bottom panel shows a_0 measured every 12 months. The stars are colored using the `matplotlib` colormap ‘Spectral’, going from red hues to blue hues akin to a rainbow’s color spectrum, where the red end indicates earlier times (1992), blue ends indicate recent times (2024) and the yellow hues indicate the midpoint in time (~2009!).

fluctuation of Δa_0 around zero for **astro-ph.EP**, **astro-ph.HE**, **astro-ph.IM**, and **astro-ph** in aggregate since 2009 can be observed. The initial increment for both **astro-ph.GA** and **astro-ph.SR** before transitioning into a consistent decrease is also observed. For **astro-ph.CO**, the initial decrease before switching over to a consistent increase is indicated as well. Lastly, for **astro-ph** as a whole since 1992, there is a consistent increment in a_0 , marked by positive Δa_0 , until about 2000, before Δa_0 is consistently below zero. In the grand scheme of things, we do want to study the arXiverse as a whole since its birth instead of just parts of it, so at this point, **astro-ph** in aggregate since 1992+ is the best proxy for arXiverse. Thus, the best proxy or measurement of the arXiverse constant is the slope measured from **astro-ph** as a whole since 1992+, at the time of writing. Hence, a_0 is quantified at 0.110 ± 0.002 paper/day/month. Last but not least, I deduce that the arXiverse is past its peak of expansion (evident by the decelerating increase in a_0 in the earlier years) and is now experiencing a decreasing rate of expansion. This means that the arXiverse, if nothing is done to alter the rate of expansion, could slow down to a crunch. I show a definitive⁹ proof of this claim in Figure 7. Let that sink in. In Figure 8, I show an illustration of what I think arXiv might be feeling upon hearing about their fate.

4.2. Limitations and Future Prospects

Throughout this work, I have performed LR fitting onto the datasets, which may be inadequate for the data. This is evident when one notices that the 95% confidence intervals of the fits in both panels of Figures 1 and 2 do not cover most of the actual data points. This suggests that the LR models may be inadequate and underfit the data. It is possible that a polynomial or exponential regression may be more suitable for different data groups. Nevertheless, to ensure a significant correlation between the average submission rate and the passage of time, I compute the Pearson p -value for all eight data groups and found that except for **astro-ph.SR** (which has a p -value of 0.026), all the other data groups have extremely low p -values. Since the significance level is typically set at 0.05, I reject the null hypothesis that there is no correlation between submission rate and time. In fact, the low p -values suggest that the observed correlation is statistically significant.

This work has used **astro-ph** as a whole as a proxy for arXiverse. This choice is motivated by the preliminary nature of this study, the fact that **astro-ph** has been there almost since the birth of the arXiverse, and

⁹ Questionable.

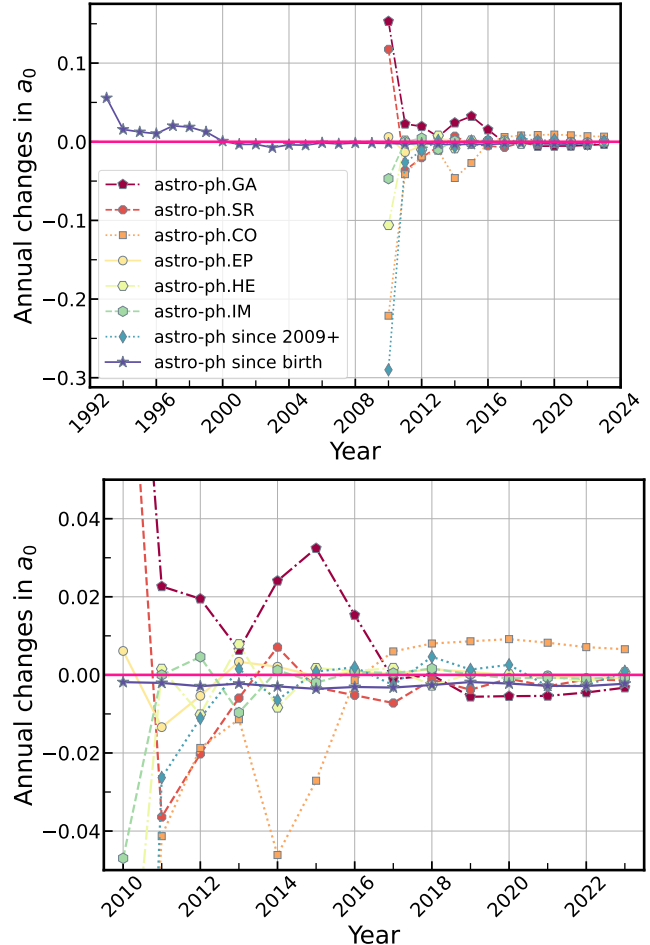


Figure 6. Annual changes in the arXiverse constant, Δa_0 , for all 8 data groups. The top panel shows the full figure, running from 1992 to date. The bottom panel shows a zoomed-in version of the top panel, focusing on the changes in a_0 from 2009 onwards and at a smaller range of the y-axis. All 8 data groups show relatively larger fluctuations in a_0 at earlier times, which then tapers off in recent times. This suggests that the change in expansion rate of the arXiverse is past its peak activity and is slowing.

that **astro-ph** has been one of the most popular categories on arXiv. However, this is not the best decision since arXiv has many more categories, even if **astro-ph** is one of the oldest or most popular categories. In the future, I propose to measure the a_0 using submission data from various arXiv categories and their respective subcategories to obtain a more robust measurement of a_0 , one that is more representative of the entire arXiverse.

5. CONCLUSION

In this work, I study the submission statistics of arXiv **astro-ph** from 1992 to date. I fit linear regression models to the time series data of the monthly submission rate of **astro-ph** as a whole and for its six subcategories indi-




Figure 7. The value of the arXiverse constant, a_0 , measured for the aggregate **astro-ph** since 1992+.

study: "the rate of expansion of the arXiverse is not constant"



Figure 8. How arXiv presumably feels about their fate.

vidually (§2). The slope of the LR fits, which tells us the monthly rate of change of submission rate, serves as the proxy for the rate of expansion of the arXiverse. I coin the term “the arXiverse constant”, a_0 , to easier quantify the rate of expansion of the arXiverse (§3). From the measured slopes as proxies for a_0 , I find that **astro-ph** as a whole (both since birth or since its split into six sub-categories), **astro-ph.GA**, **astro-ph.EP**, **astro-ph.HE**, and **astro-ph.IM** have positive a_0 — their respective submission rate (or rate of expansion) increases over time. On the other hand, **astro-ph.CO** has negative a_0 while **astro-ph.SR** has near-zero a_0 . This suggests that their submission rates decrease over time and stay constant over time, respectively. The very fact that a_0 is different across subcategories (and data groups) suggests that the rates of expansion of the (sub-)arXiverses are different from each other. This indicates that there is

a tension in the arXiverse constant, and that a_0 changes over time.

To further support the inference that a_0 is not really a constant in time, I measure a_0 across time for all the data groups (§4). The measured a_0 over time shows that it indeed changes at every given point of time. Four of the data groups show indications that they might be settling into a “Goldilocks Zone” of a_0 since the measured a_0 have been hovering right around the final measurement of a_0 . Three of the data groups indicate that there is an initial increment in a_0 before a consistent decrease in recent times; one data group shows the opposite — it has an initial decrement in a_0 before slowly increasing in recent times. I then calculate the year-to-year difference in a_0 , Δa_0 , and find the same trends in a_0 as described earlier in this paragraph. Nevertheless, there is more work to be done in the same spirit. I aim to implement various regression models onto the submission statistics and expand the dataset used for the regression modeling. I believe these two implementations will increase the robustness of the measured a_0 .

In conclusion, there is an evident tension in the arXiverse constant measured across various datasets. This is marked by the different measurements of a_0 depending on the dataset and temporal baseline used. Having said that, the best proxy for the arXiverse constant that I have at the moment is the a_0 measured from **astro-ph** as a whole since 1992+. This is because it measures the rate of expansion of **astro-ph** since its birth (and also the birth of arXiv). Thus, in March 2024, a_0 is measured to be at 0.110 ± 0.002 paper/day/month. With that said, a_0 has been consistently decreasing over time; this decrease in rate of expansion indicates that the arXiverse could slow down to a crunch. It is vital for us to play our part in ensuring that the arXiverse does not suffer this painful fate. Collectively, we can reverse this by yeeting more papers to arXiv. We all hold the power in our pens and papers to decide arXiverse’s fate. The fate is written in your hands (and perhaps the stars ★ too!). Figure 9 shows a call to action for this purpose. Remember, the pen (and paper) are mightier than the sword.

ACKNOWLEDGEMENTS

♥: The author thanks Irina Ene for the most wonderful and memeingful discussions that made this research so enjoyable and exciting. The author also thanks peer reviewers, Hitesh Kishore Das and Anshuman Acharya, for their kind efforts in reviewing this paper and ensuring the science here is logical and sound. They will be beer-ly rewarded for their contributions. “*Do a peer review, get a beer reward,*” a wise person once said. This research has made use of arXiv API, and the author




Figure 9. Call to action.

would like to thank arXiv for the use of its open-access interoperability. This research has also made use of

Software: Matplotlib (Hunter 2007), SciPy (Virtanen et al. 2020), NumPy (Harris et al. 2020) and scikit-learn (Pedregosa et al. 2011) .

REFERENCES

Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Nature*, 585, 357, doi: [10.1038/s41586-020-2649-2](https://doi.org/10.1038/s41586-020-2649-2)

Hunter, J. D. 2007, *Computing in Science & Engineering*, 9, 90, doi: [10.1109/MCSE.2007.55](https://doi.org/10.1109/MCSE.2007.55)

Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, *Journal of Machine Learning Research*, 12, 2825

Sadjadi, M. 2017, arxivscraper, doi: [10.5281/zenodo.889853](https://doi.org/10.5281/zenodo.889853)

Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *Nature Methods*, 17, 261, doi: [10.1038/s41592-019-0686-2](https://doi.org/10.1038/s41592-019-0686-2)