

Creative Commons Attribution 4.0 International (CC BY 4.0)

<https://creativecommons.org/licenses/by/4.0/>

Access to this work was provided by the University of Maryland, Baltimore County (UMBC) ScholarWorks@UMBC digital repository on the Maryland Shared Open Access (MD-SOAR) platform.

**Please provide feedback**

Please support the ScholarWorks@UMBC repository by emailing [scholarworks-group@umbc.edu](mailto:scholarworks-group@umbc.edu) and telling us what having access to this work means to you and why it's important to you. Thank you.

## PACMan2: Next Steps in Proposal Review Management

LOUIS-GREGORY STROLGER,<sup>1,2,3</sup> JAMILA PEGUES,<sup>1,\*</sup> TEGAN KING,<sup>4</sup> NATHAN MILES,<sup>5</sup> MICHELLE RAMSAHOYE,<sup>6</sup>  
KEITH CERUTI II,<sup>7</sup> BRETT BLACKER,<sup>1</sup> AND I. NEILL REID<sup>1</sup>

<sup>1</sup>*Space Telescope Science Institute, Baltimore, MD 21218, USA*

<sup>2</sup>*Johns Hopkins University, Baltimore, MD 21218, USA*

<sup>3</sup>*Morgan State University, Baltimore, MD 21251, USA*

<sup>4</sup>*National Center for Atmospheric Research, Boulder, CO 80305, USA*

<sup>5</sup>*Department of Earth, Planetary, and Space Sciences, University of California, Los Angeles, Los Angeles, CA 90095, USA*

<sup>6</sup>*Department of Mathematics and Statistics, University of Maryland, Baltimore County, Baltimore, MD 21250*

<sup>7</sup>*Baltimore Polytechnic Institute, Baltimore, MD 21209, USA*

### ABSTRACT

With the start of a new Great Observatories era, there is renewed concern that the demand for these forefront facilities, through proposal pressure, will exceed conventional peer-review management's capacity for ensuring an unbiased and efficient selection. There is need for new methods, strategies, and tools to facilitate those reviews. Here, we describe PACMan2, an updated tool for proposal review management that utilizes machine learning models and techniques to topically categorize proposals and reviewers, to match proposals to reviewers, and to facilitate proposal assignments, mitigating some conflicts of interest. We find that the classifier has cross-validation accuracy of  $80.0 \pm 2.2\%$  on proposals for time on the *Hubble Space Telescope* and the *James Webb Space Telescope*.

*Keywords:* telescopes, space telescopes, astronomers, astrophysicists

### 1. INTRODUCTION

PACMan (Strolger et al. 2017) was initially conceived as a means to reduce subjective bias and improve efficiency in the recruitment of appropriate reviewers for proposals to the *Hubble* and *James Webb* space telescopes (*HST* and *JWST*, respectively). The Space Telescope Science Institute (STScI) conducts its peer reviews of proposals for telescope time in a process generally similar to those used worldwide, for both space and ground-based observatories. Proposals are solicited through Calls for Proposals (CfPs), released annually as Cycles for each mission. Those proposals are evaluated on their scientific merit by a Time Allocation Committee (TAC), a panel of peers recruited for their expertise in broad topical science areas. On average, over 1000 proposals are received in response to each CfP, of which only the top  $\sim 20\%$  are selected. The Science Policies Group (SPG) at STScI arranges these reviews, recruiting a few hundred experts from the international scientific community in a six-month process that culminates in a set of ranked recommendations to the STScI Director.

As anticipated, *JWST* has essentially doubled the number of proposals received each year, greatly increasing the pressure in arranging thorough and unbiased reviews. There is just a finite number of qualified and available reviewers in the community, and that pool is further limited by our institutional awareness of who they are and what they do. Moreover, the global COVID pandemic has generally further reduced reviewer availability worldwide. In anticipation of a shortfall, with Strolger et al. (2017) we began to address how machine-learning tools would help streamline the review process, and to do so without increasing the potential for biases typically encountered in pressured reviews. The work was focused on automating and improving reviewer matching, and thereby expanding our reviewer pools, but still working within our peer-review framework. We also hoped not to impart subjectivity worse than as seen for typical expert reviews (Patat 2018). It does not consider changes to the way in which the peer reviews are done, e.g.,

like distributed peer-reviews (Kerzendorf et al. 2020) similar to those employed at Gemini and ESO for some parts of their biannual reviews, as there are shared concerns on strategic manipulation (Naghizadeh & Liu 2013).

Here we present the latest incarnation of the PACMan (version 2) tool-kit (King & Strolger 2023), providing a more direct proposal-to-reviewer matching, and supporting tools to help manage the reviews. PACMan2 has a set of operations to assist in the creation of panel reviews and proposal assignments. In this paper we document its core functions: topically classifying proposals and reviewers, identifying reviewers whom are best suited to reviewing specific proposal, determining close collaborators and potential conflicts of interest, and identifying resubmissions.

As with the previous version, PACMan2 uses Naive Bayesian techniques. Open-sourced python-based machine-learning toolkits such as `scikit-learn` (Pedregosa et al. 2011) have become more prevalent in the community, providing robust tools for doing all sorts of predictive analysis and classification. These tools are known to work particularly well when combined with Natural Language Processing techniques such as word-count vectorization (Kerzendorf 2019), for which toolkits have also become more widespread (e.g., `spaCy`, Honnibal et al. 2020). PACMan now utilizes the `scikit-learn` naive Bayes algorithm for multinomially distributed data, `MultinomialNB`, and takes advantage of the `spaCy` tools for lemmatization, tokenization, and word vectorization, with stop words.

## 2. IMPROVED PROPOSAL AND REVIEWER CATEGORIZATION


PACMan’s classifications are made with a multinomial naive Bayesian classifier, in a manner similar to Strolger et al. (2017) and Kerzendorf (2019), using word vectors with Term Frequency-Inverse Document Frequency (TF-IDF) weighting statistics. In brief, the Bayes’ method determines the probability of a corpus (or proposal) belonging to a pool or topical category, from the joint probability of words used in that corpus belonging to those same categories from frequency of use in training sets. TF-IDF down-weights words commonly used across multiple categories.

We classify proposals into broad pools based on science topics, correlating with the topical panels of the 2022 TAC. Those are Solar System Astronomy (abbreviated as Sol. in some areas of this text), Exoplanets and Exoplanet Formation (or Exopl.), Stellar Physics and Stellar Types (Stars), Stellar Populations and the Interstellar Medium (StPops), the Intergalactic Medium and Circumgalactic Medium (IGM), Supermassive Black Holes and Active Galaxies (AGN), Galaxies (Gal.), and Large Scale Structure of the Universe (LSS). These terms were chosen largely due to their adherence with the Unified Astronomy Thesaurus (UAT, Frey & Accomazzi 2018), which could allow for a simple, hierarchical connection of related sub-topics, or from one topical categorization to the next. This is put into practice in the recruitment of panelists who select from options of Scientific Categories and Keywords, similar to Table 3 of Frey & Accomazzi (2018), to narrow in (at a finer level) on their areas of expertise, to more directly match by hand to a similar keywords and categorization selections provided by proposers. We have only modestly tested categorizing with PACMan at such a refined level, the results of which, as shown in Figure 5 of Strolger et al. (2017), tends to show a much higher misclassification (or failure) rate that does not significantly improve with the size of the training pool.


Our classifier model is built on training from 8868 proposals received from the community in *HST* Cycles 23 through 30, for the regular CfP and for mid-cycle opportunities, in the years 2015 to 2022. The proposer-selected science categories of those proposals are shown in Figure 1, where proposals from earlier cycles were re-mapped to fit the science topics above.

### 2.1. Cross Validation

We tested our classifier model via a cross-validation method, using a *leave-one-out approach* where we specify 7 stratified K-folds of testing data, where in machine-learning parlance,  $k$  indicates the number of samples for each test (i.e.,  $k = 7$ ). Here, the data is randomly split into a training set (of 85.7% of the proposals) with stratified equal numbers of training proposals in each of the 8 categories, from which a temporary model is then derived. That model is then evaluated and scored on the remaining 14.3% of the sample (i.e., the test set). The process is repeated 7 times, making new stratified splits of training and testing along a 6:1 ratio, discarding the temporary model while retaining the score for the aggregate summary. The cross-validation accuracy of the classifier from this testing is  $80.0 \pm 2.2\%$ , while the precision and f1-measure (the harmonic mean of the precision, the fraction of categorized corpora that were accurate, and recall, the fraction of relevant corpora were accurately categorized) are  $78.8 \pm 2.7\%$  and  $78.2 \pm 2.4\%$ , respectively. The normalized confusion matrix, shown in Figure 2, emphasizes where mis-categorizations occur, in closely related topical categories. For example, Large Scale Structure of the Universe shows the largest misclassification rate (39%), with the most frequent confusion being with proposals in the Galaxies category (21% of the time). We attempt to capture the efficacy of the ‘top two’ classifications, in a method similar to Figure 3 of Strolger et al. (2017), in




**Figure 1.** Topical classifications of the proposals in the training set. Proposals were initially categorized by the proposers according to the topical panels by which they desired the proposals to be reviewed. Proposals prior to Cycle 30 required minor re-categorization (by the authors of this manuscript) to adhere to the categories used in Cycle 30.



**Figure 2.** Normalized confusion matrix of the classification model, determined from cross-validation testing. Some topical categories, e.g., Large Scale Structure of the Universe, could be confused with closely related topics, such as Galaxies.

Figure 3. Here, the frequency in which PACMan matches the manual classifications are shown in green, and how often PACMan’s second choice matches the manual classifications is shown in yellow. The ‘misclassifications’, when the manual classification is not within PACMan’s top two choices, are shown in red.



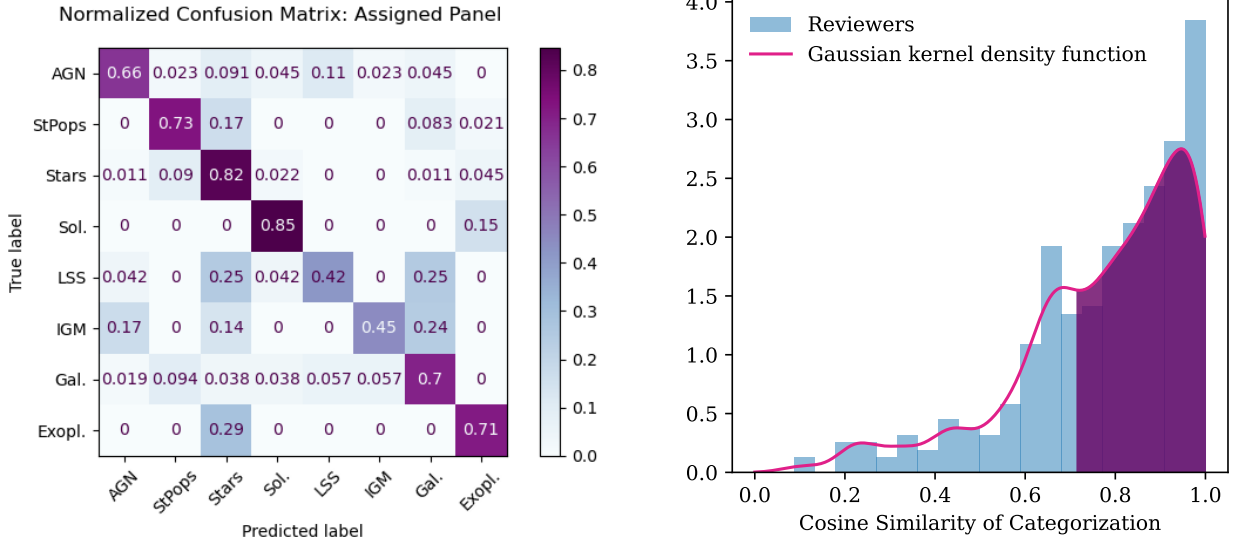
**Figure 3.** Similar to Figure 3 of Strolger et al. (2017), a comparison (or precision) of the PACMan categorization to the manual sorting of proposals from *HST* Cycles 23–30. The percentage of proposals that are similarly categorized by both are given in green. The yellow indicates how often PACMan’s second choice is the same as the ‘by hand’ classifications in that category. The red indicates the frequency of ‘misclassifications’, when the manual classification is not within PACMan’s top two choices.

As before, it is not surprising that PACMan does best in categories in which there are word tokens which unambiguously belong to said categories, such as planet names in Solar System Astronomy, or “coronagraphy” in Exoplanets and Exoplanet Formation. Topics which give PACMan the most difficulty, i.e., Stellar Populations, Galaxies, and Large Scale Structure of the Universe, have also been areas of great ambiguity to proposers. Nonetheless, PACMan has an average ‘top category’ accuracy of 84%, and a ‘top 2 categories’ average accuracy of 96%. PACMan does no worse than about a 10% classification error in any of the 8 topical categories presented. As a note, while `scikit-learn` provides a naive Bayesian method for imbalanced training sets, which may better handle our data as implied by Figure 1, our testing of the Complement Naive Bayes classification shows very little (but slightly worse) difference in overall classification performance, with a cross-validation accuracy of  $76.9 \pm 2.7\%$ , and little change in the confusion matrix.

## 2.2. Characterizing Reviewer Domain Knowledge

Until relatively recently, recruiting reviewers for the TAC has relied solely on the SPG’s familiarity with the field, and knowledge of the researchers who are active within it. A few years ago the Proposal/Person (or ProPer) Application<sup>1</sup> and its associated database was developed to maintain contact information for people who have applied for time with *HST* or *JWST*, or those associated with the STScI-organized service committees, largely for *HST* and now for *JWST*. Since then, the interface has been used internally to keep a record of self-declared areas of scientific expertise. ProPer has very much facilitated the search for panelists to serve on the TAC topical panels, and through the addition of UAT terms as keywords, matching proposals to reviewers has also been made easier (more on that in Section 3). Yet, this extensive record only exists for those who have applied for time with these observatories, or in one way or another, connected to STScI service committees. It does not reach to, for e.g., primarily ground-based astronomers, or astronomers focusing on other wavelengths, or in many cases, theorists, instrumentalists, or other non-observers. These investigators are likely equally knowledgeable about the science to be evaluated, without being experts on the

<sup>1</sup> ProPer can be found on <https://proper.stsci.edu/proper/authentication/>.



**Figure 4.** Comparisons of reviewer categorizations via PACMan (predicted label) to those provided by the reviewers themselves (true label). *Left:* the normalized confusion matrix between reviewer-supplied categories of expertise and PACMan category vectors on reviewers. *Right:* normalized distribution of cosine similarity between category vectors. The dark purple region marks the top 68<sup>th</sup> percentile of the gaussian kernel distribution (purple line).

use of these facilities. We wish to apply our trained model to categorize potential reviewers under the same pools as proposals, with a goal of facilitating and expanding potential reviewer pools for each of the appropriate topical review panels.

Assuming a record of domain knowledge exists in an author’s bibliography, we create evaluation corpora from Astrophysics Data System (ADS)<sup>2</sup> queries for the last 10 years worth of refereed paper abstracts, and in lieu of proposals, from a list of proposed referees. These authors are identified in queries to ADS (made via the API<sup>3</sup>) by first name and last name. It would be ideal to use a persistent identifier (e.g., ORCID<sup>4</sup> iDs) that distinguishes authors with the same or similar names, and unambiguously identifies those who may have changed their names over their careers. But unfortunately, from our experience, the record of ORCID iDs has been incomplete for many authors in our field, and therefore an unreliable query to ADS. The occasional ambiguous query from names alone results in more records attributed to a given name than are warranted, and a broader characterization (and more frequent proposal matching, see Section 3 below) than would be expected. At the moment, there is little to guard against this, and therefore these results (as with all PACMan results) should be used as recommendations to guide to reviewer-panel organizers and managers. We remain hopeful that the usage of such identifiers expands greatly, and envision incorporating the ORCID public API with our code in a future release.

PACMan interprets these corpora, applying the same lemmatization, tokenization, and word vectorization as is done with proposals, and the same trained classification. As an illustration of the effectiveness of this method, Figure 4 compares reviewer categorizations from PACMan to those supplied by the reviewers themselves for *HST* Cycle 30. Here, the internal and external reviewers provide their expertise in ProPer through selections of sub-category keywords.

As a test of the PACMan-assigned domain knowledge, against user-supplied information, we make use of the fact that it is often the case that reviewers have expressed at least some expertise in a wide-range of keywords that map-back to more than one topical category. Moreover, PACMan categorizations are actually probabilities of expertise in each topical category. Both tidbits of information can be used and compared, converting the frequency of times certain categories are selected from mapped-back keywords into a mock category vector ( $\vec{v}_1$ ), and using the PACMan probabilities as a separate mock category vector ( $\vec{v}_2$ ), we can compute the cosine similarity of these two mock category

<sup>2</sup> <https://ui.adsabs.harvard.edu>.

<sup>3</sup> <https://ads.readthedocs.io>.

<sup>4</sup> <https://orcid.org>. The public API can be found on <https://info.orcid.org/documentation/features/public-api/>.

vectors as:

$$S_C(\vec{v}_1, \vec{v}_2) = \frac{\vec{v}_1 \cdot \vec{v}_2}{\|\vec{v}_1\| \|\vec{v}_2\|}. \quad (1)$$

The right panel of Figure 4 shows the cosine similarity ( $S_C$ ) distribution, which indicate that the PACMan categorizations agree, for the most part, with the reviewer-supplied categorizations. In finer detail, these 8-component  $S_C$ 's are largely very high (where  $S_C \equiv 1$  is a perfect match), with the highest 68% at  $S_C > 0.7$ . These mock 8-component vectors were only constructed for this test of domain knowledge, and are not used elsewhere in the PACMan code-base. Instead, the multi-component word vectors, as further described in Section 3, are used and are naturally more complicated, and as such result significantly lower (i.e.,  $S_C \ll 1$ ) values.

It's not always straight-forward for an individual to characterize their expertise in limited topical categories. Further, there are known biases in how individuals characterize the scope of their work. PACMan potentially captures a better, less-biased characterization of that expertise from the publications these individuals author and co-authors. The interesting collateral outcome of these model-based characterizations may be the result of a wider range of expertise, for a given reviewer, than said reviewer may have provided themselves.

### 3. REVIEWER-TO-PROPOSAL MATCHING

The categorization tool provides a base method for identifying which topical panel proposals should be sent to, and who would be adroit to serve on those review panels. The next step is direct proposal-to-reviewer assignments. The *HST* reviews previously made use of the user-defined categories and keywords, and assignments were primarily made based on the number of proposal-reviewer overlapping keywords. Other considerations, such as load balance and reducing the conflicts of interests (discussed in Section 4) would further solidify those assignments. However, experience has shown that user-provided categorizations and keywords can be misused, misconstrued, or underutilized, leading to inaccuracies in those assignments.

There are well-established methods for pairwise comparisons of corpora, determining likenesses (or affinities), or differences. One simple evaluator is the cosine similarity ( $S_C$ ), as described in Equation 1. Forgoing the categorizations described in Section 2, and the mock 8-component vectors from the Section 2.2, this method directly compares the TF-IDF word-frequency multi-component vectors to each other to determine which corpora are best aligned with a given reviewer's ADS bibliography (as described in Section 2.2). While  $S_C$  scores closest to a value of 1 represent the most ideal matches, it would be nearly impossible for a proposal to use *exactly* the same words, or in the same frequency, as those used in the manuscripts of an author's bibliography. It begs the question of what range of  $S_C$  values which represent scores for 'experts', or those best qualified to review a given proposal?


To test this we sought to compare the list of potential reviewers generated through ProPer for the *JWST* Cycle 1 TAC review, against the proposals received for that CfP deadline in 2020. We constructed a few tests of this list of reviewers against the proposals, the first evaluating all 185 potential reviewers against all 1171 proposals, totaling 216,635 scores, to create a baseline of  $S_C$  score distribution (see Figure 5). As one would expect, some subset of reviewers have the highest scores on a few proposals, with the top 15% in the range  $S_C > 0.046$ . But most comparisons would be average to poor, with  $\bar{S}_C = 0.03^{+0.02}_{-0.01}$  and  $\text{Mo}(S_C) = 0.02$ . To better isolate the range best scores, we looked more closely at just those top-scoring reviewers for each proposal, then the highest  $S_C$ , and then the top six reviewers assigned to each proposal<sup>5</sup>. The distribution of those  $S_C$ 's for the top-six reviewers was clearly distinguished from the baseline, with  $\bar{S}_{(C, \uparrow 6)} = 0.08^{+0.03}_{-0.02}$  and  $\text{Mo}(S_{C, \uparrow 6}) = 0.07$ . And the best reviewers tended to even higher scores, with  $\bar{S}_{(C, \text{Best})} = 0.10^{+0.04}_{-0.03}$  and  $\text{Mo}(S_{C, \text{Best}}) = 0.08$ . These results are also shown in Table 1. From the distributions, it seems  $S_C \geq 0.06$  generally denotes a range in which a reviewer is sufficiently expert enough to review a given proposal, and that assignments of  $S_C < 0.05$  should be done with some caution.

It is not always the case that the proposers themselves are the best suited to evaluate the subject matter of their proposals. Figure 5 also shows the distributions of  $S_C$  scores for the proposal principal investigators (PIs) against their own proposals, showing median expertise on par with the best six reviewers, but with a wider distribution, with  $\bar{S}_{(C, \text{PIs})} = 0.08^{+0.06}_{-0.04}$ . So, while it encompasses the range of some of the top-most reviewers, it also covers some of the range of the least expert of reviewers.

This exemplifies some issues in how  $S_C$  is determined, and perhaps serves as a caution for using these values wholesale without due consideration. It is very possible for any author to have expertise in several areas, and to develop expertise

<sup>5</sup> Six was somewhat arbitrarily chosen. Proposals are generally assigned to three readers, which was doubled to guard against conflicts of interest which could disqualify some readers.





**Figure 5.** Distributions of proposal-to-reviewer cosine similarity ( $S_C$ ) scores are shown, along with their gaussian kernel distributions, for the subgroups of reviewers selected by expertise shown in the legend and in Table 1. Vertical lines indicate the median  $S_C$  values, and shaded regions show the 68% confidence intervals around those medians. The number of within those subgroups are also indicated in the legend.

	$\bar{S}_C$	Mo( $S_C$ )
Any Reviewer	$0.03^{+0.02}_{-0.01}$	0.02
Best Six Reviewers ( $\uparrow 6$ )	$0.08^{+0.03}_{-0.02}$	0.07
Top Reviewers (Best)	$0.10^{+0.04}_{-0.03}$	0.08
PIs as Reviewers (PIs)	$0.08^{+0.06}_{-0.04}$	0.05

**Table 1.** Medians, 68%-confidence regions, and modes for the subgroups of reviewers selected by expertise, in the evaluation of range of validity of proposal-to-reviewer cosine similarities ( $S_C$ ).

in emergent areas. It is this latter case that PACMan would be particularly poor at, determining their overall expertise on a topic for which they haven’t written much about. PACMan can, however, respond with who amongst the pool of names given is best to review a given proposal, which would be a proper interpretation of its results.


#### 4. IDENTIFYING CLOSE-COLLABORATORS AS POTENTIAL CONFLICTS OF INTEREST

The categorization models and  $S_C$  tools described above provide a great starting point for recruiting panels and making reader assignments. However, such assignments should also be weighed against other considerations, like assignment load, and more importantly, potential for conflicts of interest. Generally, conflicts of interest arise when a reviewer’s interest, in direct or tractable resource (or financial) gain, is at odds with their ability to provide an unbiased review of a given proposal. This definition can be expanded to cover even the perception of conflict of interest, to an outside observer, even though the actual gain to the reviewer may be minimal.

The SPG at STScI has taken great strides to reduce such biases, implementing a dual-anonymous peer review<sup>6</sup> in which the identities of the proposers are not directly known to the readers as they do their assessment of scientific merit. On the one hand, it would seem that such a policy addressing identity bias would have the collateral benefit of essentially eliminating conflicts of interest, *ab initio*— if the reviewer does not know the identities of the proposers, how would they know if they stand to gain from its award or rejection? But on the other hand, it is important to note that the goal of those dual-anonymous policies are not to completely erase all identity of the proposers, rather to

<sup>6</sup> See <https://outerspace.stsci.edu/display/APRWG>.





**Figure 6.** The number of proposal pairs with  $N$  alike winnowed 7-grams (w7g) from Cycle 29, compared against Cycles 23 through 29. From manual reviews, proposal pairs that share more than one-hundred w7g in common are essentially resubmissions.

obscure them enough so that those identities are not a focus of discussion during the review of scientific merit. There may yet be room for conflicts of interest to arise. The unfortunate side effect of the process is that readers cannot simply review the authorship and directly know if they have a conflict of interest to declare. The onus of determining these conflicts defaults back to the SPG and requires tools to help facilitate.


The queries to the ADS interface used in panel assignments and reviewer assignments also provide a means for identifying ‘close-collaborators’, at least in published works. From these queries, PACMan provides lists of close collaborators with a given reader on three escalating levels: any co-authors, co-authors on 3 or more papers, and co-author as a primary author (any of the first three names) in the past 3 years. These criteria provide a range of minor, moderate, and major conflicts of interests, respectively, which are appropriately weighed with PACMan-recommended panel and reviewer assignments by the SPG in finalizing actual assignments.

## 5. DETERMINING RESUBMISSIONS

It is often very useful to know what proposals are re-worked resubmissions from previous cycles, and without an encyclopedic knowledge of every proposal ever received, those can be very hard to identify. PACMan provides a tool to determine what proposals in past submissions have close similarity to those submitted for the current cycle. Those proposal duplications are determined via hashed winnowed  $k$ -gram similarity, using the method of Citron & Ginsparg (2015), developed for [arXiv.org](https://arxiv.org). Using  $k$ -grams, or sentences of greater than  $k$ -words length, the method essentially creates a fingerprint of each proposal, by collecting the set of  $k$ -grams within each. A proposal of  $n$  words has  $n - (k + 1)$   $k$ -grams. Using **Fingerprint**<sup>7</sup>, each  $k$ -gram is hashed into a unique long integer identifier by using the remainder of dividing the  $k$ -gram by `sys.maxsize`<sup>8</sup>, with the string of hashes forming a corpus’s “fingerprint”, easily indexed in a computer’s RAM. One can further facilitate calculations by taking advantage of the fact that  $k$ -grams have overlap, and therefore redundant information that unnecessarily increases the size of these fingerprints. The hashes are thus winnowed by a factor of  $t > k$ , where each proposal is composed of  $n - t + 1$  windows, each with  $t - k + 1$  winnowed  $k$ -grams. Adopting the lowest hash in rolling windows of length  $t$  further compresses the size of corpora fingerprints. Based on Citron & Ginsparg (2015) recommendations, we use  $k = 7$  and  $t = 12$ , compressing fingerprints of winnowed 7-grams (w7g’s) effectively by factors of 3 to 4.

<sup>7</sup> **Fingerprint** codebase: <https://github.com/kailashbuki/fingerprint>. under an MIT license.

<sup>8</sup> The maximum size object containers (lists, strings, etc.) can have in python, which for 64-bit machines is  $2^{63} - 1 = 9,223,372,036,854,775,807$ . The method is discussed in great detail in the **Footprint** GitHub repository.



**Figure 7.** The number of suspected resubmissions ( $w7g > 100$ ) in the last six cycles, shown as a stack histogram from within the same year (from mid-cycle proposals, in green), then including resubmissions from the last cycle (in orange), then from  $\leq 3$  years (in blue).

Figure 6 shows the distribution of  $w7g$  values for the 1130 proposals received in Cycle 29, in comparison to the 7805 proposals in Cycles 23 through 29. Similar distributions were drawn for each Cycle, 23 through 30, and reviewed manually, from which it was determined that  $w7g \gtrsim 100$  were highly likely resubmissions. From these manual reviews, it was also noted that the number of resubmissions appeared to be growing. Adopting  $w7g > 100$  as the cutoff, we show (in Figure 7) that the number of suspected resubmissions has gone up in the last 5 years, most notably in Cycle 30. Most of the resubmissions tend to be from just the previous cycle ( $\approx 80\%$ ), with most of the remainder coming from within the last three cycles ( $\approx 14\%$ ). The subsequent remainder of  $\approx 6\%$  are resubmissions of mid-cycle proposals.

## 6. SUMMARY

PACMan provide a useful toolkit for proposal and reviewer management in the STScI peer-review system, and similarly should have broad applicability to other peer-review management groups. The PACMan version 2 codebase is available on zenodo (King & Strolger 2023), and maintained on github, <https://github.com/spacetelescope/PACMan>.

## ACKNOWLEDGMENTS

We thank our anonymous referee for important improvements to this manuscript. We also thank Alessandra Alossi, Claus Leitherer, Christine Chen, Kathrine Alatalo, Molly Peebles, and Laura Watkins for their discussions in preparing and implementing PACMan. MR thanks the STScI Space Astronomy Summer Program. KC thanks The Baltimore Polytechnic Ingenuity Project for enabling their participation in this work. This research has made use of NASA’s Astrophysics Data System Bibliographic Services, in particular, the API for the Astrophysics Data System.

*Author contributions:* We use the CRT standard (<https://casrai.org/credit/>) for reporting author contributions. Conceptualization: L.G.S. Data curation: L.G.S., B.B., I.N.R. Formal analysis: L.G.S., J.P., K.C. Investigation: L.G.S., J.P., T.K., N.M., M.R., K.C., B.B., I.N.R. Methodology: L.G.S. Software: T.K., N.M., M.R. Supervision: L.G.S. Validation: L.G.S. Writing—original draft: L.G.S. Writing—review & editing: L.G.S., J.P., T.K., N.M., M.R., K.C., B.B., I.N.R.

*Facilities:* Hubble Space Telescope, James Webb Space Telescope

*Software:* **Astropy** (Astropy Collaboration et al. 2013), **Matplotlib** (Hunter 2007), **Scikit-learn** (Pedregosa et al. 2011), **spaCy** (Honnibal et al. 2020) , **Fingerprint**.

## REFERENCES

- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, 558, A33, doi: [10.1051/0004-6361/201322068](https://doi.org/10.1051/0004-6361/201322068)
- Citron, D. T., & Ginsparg, P. 2015, *Proceedings of the National Academy of Sciences*, 112, 25, doi: [10.1073/pnas.1415135111](https://doi.org/10.1073/pnas.1415135111)
- Frey, K., & Accomazzi, A. 2018, *ApJS*, 236, 24, doi: [10.3847/1538-4365/aab760](https://doi.org/10.3847/1538-4365/aab760)
- Honnibal, M., Montani, I., Van Landeghem, S., & Boyd, A. 2020, *spaCy: Industrial-strength Natural Language Processing in Python*, Zenodo, doi: [10.5281/zenodo.1212303](https://doi.org/10.5281/zenodo.1212303)
- Hunter, J. D. 2007, *Computing in Science and Engineering*, 9, 90, doi: [10.1109/MCSE.2007.55](https://doi.org/10.1109/MCSE.2007.55)
- Kerzendorf, W. E. 2019, *Journal of Astrophysics and Astronomy*, 40, 23, doi: [10.1007/s12036-019-9590-5](https://doi.org/10.1007/s12036-019-9590-5)
- Kerzendorf, W. E., Patat, F., Bordelon, D., van de Ven, G., & Pritchard, T. A. 2020, *Nature Astronomy*, 4, 711, doi: [10.1038/s41550-020-1038-y](https://doi.org/10.1038/s41550-020-1038-y)
- King, T., & Strolger, L. 2023, *spacetelescope/PACMan: Public release of version 2, v2.0*, Zenodo, doi: [10.5281/zenodo.7662460](https://doi.org/10.5281/zenodo.7662460)
- Naghizadeh, P., & Liu, M. 2013, *arXiv e-prints*, arXiv:1307.6528. <https://arxiv.org/abs/1307.6528>
- Patat, F. 2018, *PASP*, 130, 084501, doi: [10.1088/1538-3873/aac463](https://doi.org/10.1088/1538-3873/aac463)
- Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, *Journal of Machine Learning Research*, 12, 2825
- Strolger, L.-G., Porter, S., Lagerstrom, J., et al. 2017, *AJ*, 153, 181, doi: [10.3847/1538-3881/aa6112](https://doi.org/10.3847/1538-3881/aa6112)