Computation and Performance Studies
Volume 1 · Number 1 · 2026

Fewer Working Pixels, Lower Foreground F1: A Measured Document-Processing Tradeoff

Gregory Franklin[1] · Ethan Hughes[2] · Paul Lewis[3]

[1] Australia · [2] New Zealand · [3] Netherlands

Original publication: October 28, 2026 | Article 666

Abstract

Reducing a document's working grid saved processing time while losing foreground accuracy in this experiment. Ten H-DIBCO 2018 pages were processed at native resolution and at approximately half each dimension, with a correspondingly reduced local threshold window. The smaller pipeline used about one quarter of the working pixels and had a geometric mean latency ratio of 0.300, while mean foreground F1 fell from 0.648 to 0.629. Timing includes reduction and binary expansion back to the original grid. The result describes a specific digital processing policy on one CPU, rather than a general speed or quality guarantee.

Exploratory research draft. One public document cohort and one CPU timing session; final scientific, originality, and document review remains pending.

Fewer Working Pixels, Lower Foreground F1: A Measured Document-Processing Tradeoff

> Exploratory research draft. One public document cohort and one CPU timing session; final scientific, originality, and document review remains pending.

Abstract

Reducing a document's working grid saved processing time while losing foreground accuracy in this experiment. Ten H-DIBCO 2018 pages were processed at native resolution and at approximately half each dimension, with a correspondingly reduced local threshold window. The smaller pipeline used about one quarter of the working pixels and had a geometric mean latency ratio of 0.300, while mean foreground F1 fell from 0.648 to 0.629. Timing includes reduction and binary expansion back to the original grid. The result describes a specific digital processing policy on one CPU, rather than a general speed or quality guarantee.

1. Resolution is a resource decision with an output cost

Document-processing pipelines can reduce their working image dimensions to avoid filtering every original pixel. That can be attractive when many pages must be processed, but the smaller representation also changes what the algorithm observes. Fine variations are averaged, thresholds are computed on another grid, and the resulting mask must be mapped back to the original document. A resource comparison should therefore evaluate both the actual time saved and the accuracy of the final output on a common support.

The research question asks what foreground accuracy and measured processing-time tradeoff follows from halving the working grid while approximately rescaling a frozen local threshold window. The official H-DIBCO 2018 cohort Pratikakis et al., 2018 supplies ten handwritten pages and paired binary references. We process each page through two declared policies and score both against the full-resolution reference. This prevents a smaller output array from appearing accurate simply because its difficult fine-scale content has also been removed from the evaluation target.

Meng et al., 2025 develop zero-shot inpainting of oracle-bone inscription rubbings with a visual Mamba component. Its restoration setting motivates careful treatment of document content when reducing computation. Our experiment is a binarization study, not an implementation of that model. It does not reconstruct absent strokes, assess archaeological authenticity, or measure character recognition. The connection is the practical need to distinguish a cheaper image-processing representation from evidence that meaningful content has been preserved.

Two propositions guide the comparison. A smaller grid should reduce some numerical filtering work, but actual latency also depends on resampling, allocation, and expansion overhead. The change in foreground accuracy is an empirical question because averaging can remove noise as well as useful detail. We therefore do not assume that the smaller policy must improve or worsen every page. The design saves both page-level accuracy differences and repeated measured processing times, with neither outcome used to retune the policies on held labels.

2. Digital reduction and the information boundary

Sampling theory Shannon, 1949 gives a conceptual reason to distinguish reducing a representation from preserving all its information. Its idealized signal assumptions are not established for these already sampled document images. We do not claim a band-limited recovery guarantee or calculate an optical sampling limit. The relevant empirical fact is that digital area reduction combines original pixels before the threshold rule acts. Later nearest-neighbor expansion enlarges the binary mask but does not restore the original intensity observations.

Structured recovery work Candes and Plan, 2010 likewise depends on assumptions about the data and observation process. No such recovery method is used here. The smaller pipeline simply produces a mask from its reduced grayscale image. It may retain sufficient information for some foreground decisions and lose information important to others. That limitation should be evaluated through the actual target mask, rather than treating a reduction in array size as equivalent to lossless compression or a reconstruction guarantee.

The intervention also differs from acquiring the same document with another camera. Unprocessing research Brooks et al., 2018 and low-light reconstruction work Chen et al., 2018 concern imaging pipelines and raw measurement conditions that our released images do not identify. We have no camera calibration, optical resolution, or physical pixel size for this comparison. Half-resolution therefore means a software transformation of the available array, not a verified change in sensor resolution, scanner quality, or capture distance.

The image algorithm is Sauvola's local binarization rule Sauvola and Pietikäinen, 2000. Its native configuration uses a fifty-one-pixel window and coefficient 0.2, inherited from an earlier-cohort source selection record. The half-grid configuration uses a twenty-five-pixel window and the same coefficient. Twenty-five approximates half of fifty-one while remaining an odd window size. It does not exactly preserve an independently calibrated physical neighborhood. Both grid formation and this window rounding are components of the processing-policy change.

Otsu's global histogram threshold Otsu, 1979 provides another binarization construction and is the comparator in the companion cohort study. We do not switch between global and local algorithms in the present timing contrast. Keeping the algorithmic family fixed focuses the question on its working representation and associated window adjustment. A separate reduced-grid Otsu experiment could have another quality-time relationship, but that result is not available from the local-rule comparison reported here.

3. Two complete processing paths

The source archive contains ten native-resolution images and their paired binary ground truths. ZIP integrity, image-reference dimensions, binary reference values, source URLs, and file hashes are checked and retained. Grayscale source arrays are decoded before timing. The native pipeline converts the array to floating numerical values, computes local mean and variance with reflection boundaries, applies the declared threshold, and returns a binary mask on the original grid. Its reference standard-deviation scale remains 128 for the eight-bit intensity range.

The reduced pipeline uses Pillow's BOX reduction to a width and height equal to the ceilings of half the originals. It then applies the twenty-five-pixel threshold window to that smaller eight-bit grayscale array. Finally, nearest-neighbor expansion maps the binary result back to the exact original dimensions. The output is scored at that size. Because original dimensions can be odd, the total working-pixel ratio is not exactly one quarter; the observed ratio across the cohort is 0.2502058.

The reduction combines area averaging, eight-bit representation of the smaller image, a new neighborhood grid, and binary expansion. We do not attribute any accuracy change solely to one of those operations. A crossed experiment holding the intermediate intensity precision or window footprint separately fixed would be required for that decomposition. Here the policy is a concrete deployable sequence of operations, which makes its total quality-time tradeoff measurable without claiming to isolate every internal mechanism.

Numerical array and filtering operations use NumPy Harris et al., 2020 and SciPy Virtanen et al., 2020. Local variance is computed from the second moment minus the squared mean, with negative roundoff clamped to zero before the square root. Numerical-analysis discussion Higham, 2002 and floating-point arithmetic guidance Goldberg, 1991 motivate recording that detail. The clamp is part of the implemented algorithm, not a data exclusion or a mechanism for improving a selected page after examining its reference mask.

The native baseline intentionally matches the companion earlier-to-later cohort analysis. All ten native per-page F1 values agree with that saved record within the declared numerical tolerance. This is a consistency check on shared inputs and implementation. It is not an independent replication in another dataset. The current question changes the processing grid and records timing; the companion question selected a rule on one competition cohort and evaluated it on another.

4. What the timer includes

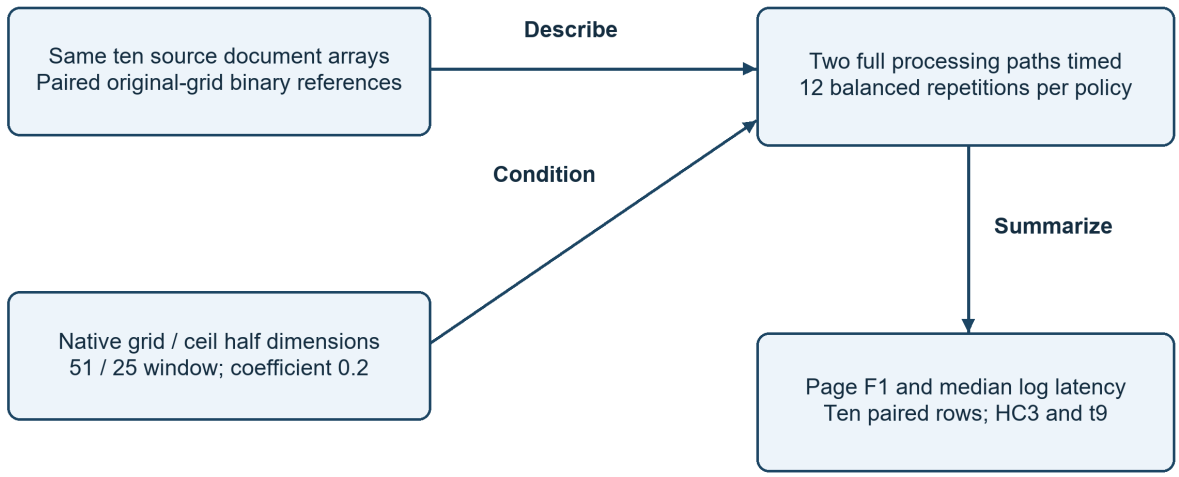


Figure 1. Shared source pages, native and reduced processing paths, and original-grid outcomes.

Figure 1 defines a paired processing-policy comparison. Each path begins from the same already decoded grayscale array, and each returns an original-grid binary mask. Timing includes array conversion and threshold computation for the native path. For the reduced path it additionally includes image reduction, conversion of the smaller array, and expansion of its binary output. Input decoding, disk access, reference loading, metric calculation, source parameter selection, and writing result files are outside the measured interval.

The timing record contains two warm-up calls per policy per page, followed by twelve measured repetitions per policy. For every page, six repetitions run native first and six run reduced first; their order is shuffled with seed 20261107. The outcome for each policy is that page's median measured duration.

This reduces dependence on a single observed call but does not remove all background-load effects. No processor affinity, thread-count override, or claim of an otherwise idle machine is imposed.

The measured session uses the recorded AMD64 processor environment, Python 3.12.14, NumPy 2.5.4, SciPy 1.18.1, and Pillow 13.0.0. The monotonic performance-counter resolution is one ten-millionth of a second. These details make the scope of the numbers inspectable. Benchmarking cautions in Mytkowicz et al., 2009 and Kalibera and Jones, 2013 explain why another session or machine can yield different latency values. Reproduction should verify the processing and quality calculations without requiring timing files to be byte-identical.

The roofline framework Williams et al., 2009 provides context for distinguishing data volume from execution constraints. We do not measure memory bandwidth, processor utilization, or a hardware roofline. A quarter of the working pixels is therefore not a proof of a quarter of the time. Resampling and expansion add work, and array allocation or caching can alter the relationship. The actual paired timer measures the declared paths, while the pixel ratio describes their representation sizes.

5. Accuracy and latency estimands

Foreground F1 is calculated against every original-grid reference pixel. Black reference pixels are the foreground class. The two policies share the same reference and metric definition for each page. The quality contrast is reduced-pipeline F1 minus native-pipeline F1. This is a difference in a bounded binary-classification summary, not an intensity mean squared error, an OCR score, or a measure of whether an individual stroke remains meaningful to a human reader.

The latency contrast is the natural logarithm of the reduced page-median duration divided by the native page-median duration. Exponentiating its equal-page mean yields a geometric mean latency ratio. We also report the arithmetic mean of each policy's ten page medians. These summaries weight the cohort differently and should not be conflated. None is a measurement of whole-archive throughput that includes disk reading, queueing, or parallel execution of many documents.

Two intercept-only OLS models summarize the paired page contrasts, with HC3 standard errors and t references using nine degrees of freedom. The robust covariance tradition White, 1980 motivates stating the estimator, while dependence guidance Cameron and Miller, 2015 discourages counting pixels or timing repetitions as additional independent document observations. The regression N is ten for each outcome. No controls or fixed effects are added, and R-squared is zero because the models estimate means only.

Intervals condition on these source pages, the inherited rule, and the observed timing session. They do not incorporate a new hardware sample, a repeated source-parameter search, or population coverage of scanner and handwriting types. Structured validation Roberts et al., 2017 and model-selection concerns Varma and Simon, 2006 help define that boundary. The held binary references are used to evaluate the declared reduction policy, not to choose a more favorable reduced window after inspecting its quality loss.

6. The smaller pipeline saves time and loses F1

Native mean foreground F1 is 0.6479992; the reduced pipeline has 0.6289004. Their paired mean difference is negative 0.0190987, with standard error 0.0048985 and interval from negative 0.0301800 to negative 0.0080175. Nine pages have a negative F1 difference and one has a positive difference. This is evidence of a quality loss for the declared finite-cohort policy comparison, while leaving readability and other image-quality dimensions unmeasured.

The geometric mean latency ratio is 0.2999117. Its log-scale mean is negative 1.2042672, with standard error 0.0470388 and interval from negative 1.3106763 to negative 1.0978581. Across page medians, the arithmetic mean duration is 36.179085 milliseconds for native processing and 11.250155 milliseconds for the reduced path. These are measured RAM-resident processing durations in one session. They are not an end-to-end service-level latency or a universal hardware speedup claim.

Table 1. Paired original-grid quality and processing-latency contrasts

Model	Term	Estimate (SE)	95% CI	N
M1	Mean outcome/contrast	-0.0191 (0.0049)	[-0.0302, -0.0080]	10
M1: fit	R-squared = 0.0000	-	-	10
M2	Mean outcome/contrast	-1.2043 (0.0470)	[-1.3107, -1.0979]	10
M2: fit	R-squared = 0.0000	-	-	10

Model labels: M1 = Half-resolution minus native foreground F1; M2 = Half-resolution versus native median-processing log latency ratio.

Table 1 reports intercept-only OLS means, HC3 standard errors, t9 intervals, and ten page observations. M1 estimates reduced-minus-native foreground F1; M2 estimates the log of the reduced-to-native page-median processing-time ratio. No covariate controls or fixed effects enter either model. Timing includes each declared processing path but excludes disk access and evaluation. The quality and timing contrasts use the same ten pages and are not independent replications of one another.

The per-page quality differences are not identical in magnitude. The largest loss in this cohort is about 0.045374 on page eight, while page nine improves by about 0.004381. We do not change the reduction rule for those pages or remove them as exceptions. Influence diagnostics such as those introduced by Cook, 1977 could examine a richer fitted model, but the saved paired page record already makes the simple heterogeneity visible. The average tradeoff should retain that distributional context.

7. What kind of quality was preserved?

Foreground agreement is only one image evaluation. Structural similarity Wang et al., 2004, learned perceptual similarity Zhang et al., 2018, structure-texture similarity Ding et al., 2020, and feature similarity Lin Zhang et al., 2011 define other ways to compare images. No score from those methods is inferred from the F1 change. In particular, a smoother-looking reduced image could still lose a thin foreground element, while a detailed mask can include false foreground. Neither possibility is resolved by the current table without calculating the relevant additional outcome.

Blind assessment methods such as BRISQUE Mittal et al., 2012, NIMA Talebi and Milanfar, 2017, MUSIQ Ke et al., 2021, and MANIQA Yang et al., 2022 also target different evaluative objects. They are not run here, and their expected behavior cannot be deduced from the latency ratio. An apparently acceptable global quality score would not establish that binarized content is preserved. A future comparison should specify whether its objective is pixelwise foreground extraction, perceived appearance, character recognition, or a downstream archival decision, then measure that object directly.

The resource conclusion also remains conditional. A user with a tolerable foreground-loss budget might value the observed processing reduction; another task might treat the lost content as unacceptable. We do not choose that utility on behalf of an archive or infer it from p-values. Acquisition costs, storage formats, transfer latency, and OCR consequences are absent. The result identifies the measured quality-time pair

for one concrete policy, providing evidence for a later decision rather than declaring a cost-optimal operating point.

8. Reproducibility and conclusion

The saved record contains the native and reduced dimensions, threshold settings, per-page F1 values, all 240 measured call durations, order positions, warm-up count, environment information, and structured estimates. Original benchmark attribution and source conditions accompany the data. Documentation principles Gebru et al., 2021 and FAIR stewardship Wilkinson et al., 2016 favor retaining those construction details. Numerical quality outputs can be reproduced from the images, while latency should be treated as an observed measurement with a new session record when rerun.

This experiment finds a clear finite-cohort tradeoff: roughly one quarter of the working pixels and substantially lower measured processing time accompany a modest reduction in mean foreground F1. The reduced pipeline is a different representation and processing path, including resampling and window rounding. Reporting that full path, scoring on the original grid, and retaining page-level differences make the conclusion more useful than either a pixel-count argument or a speed figure alone.

References

1. Talebi, Hossein and Peyman Milanfar. "NIMA: Neural Image Assessment" *arXiv preprint arXiv:1709.05424*, 2017. https://arxiv.org/abs/1709.05424.
2. Kalibera, Tomas; Jones, Richard. (2013). Rigorous benchmarking in reasonable time. Proceedings of the 2013 international symposium on memory management. https://doi.org/10.1145/2464157.2464160.
3. Varma, Sudhir and Richard Simon. “Bias in error estimation when using cross-validation for model selection.” *BMC Bioinformatics*, vol. 7, no. 1, 2006. https://doi.org/10.1186/1471-2105-7-91.
4. Ke, Junjie, et al. "MUSIQ: Multi-scale Image Quality Transformer" *arXiv preprint arXiv:2108.05997*, 2021. https://arxiv.org/abs/2108.05997.
5. Lin Zhang, ; Lei Zhang, ; Xuanqin Mou, ; Zhang, D.. (2011). FSIM: A Feature Similarity Index for Image Quality Assessment. IEEE Transactions on Image Processing. https://doi.org/10.1109/tip.2011.2109730.
6. Williams, Samuel; Waterman, Andrew; Patterson, David. (2009). Roofline: an insightful visual performance model for multicore architectures. Communications of the ACM. https://doi.org/10.1145/1498765.1498785.
7. Goldberg, David. (1991). What every computer scientist should know about floating-point arithmetic. ACM Computing Surveys. https://doi.org/10.1145/103162.103163.
8. Mytkowicz, Todd; Diwan, Amer; Hauswirth, Matthias; Sweeney, Peter F.. (2009). Producing wrong data without doing anything obviously wrong!. Proceedings of the 14th international conference on Architectural support for programming languages and operating systems. https://doi.org/10.1145/1508244.1508275.
9. Sauvola, J.; Pietikäinen, M.. (2000). Adaptive document image binarization. Pattern Recognition. https://doi.org/10.1016/s0031-3203(99)00055-2.
10. White, Halbert. “A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity.” *Econometrica*, vol. 48, no. 4, 1980, pp. 817-838. https://doi.org/10.2307/1912934.
11. Gebru, Timnit, et al. "Datasheets for datasets" *Communications of the ACM*, 2021. https://doi.org/10.1145/3458723.
12. Roberts, David R., et al. “Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure.” *Ecography*, vol. 40, no. 8, 2017, pp. 913-929. https://doi.org/10.1111/ecog.02881.
13. Brooks, Tim; Mildenhall, Ben; Xue, Tianfan; Chen, Jiawen; Sharlet, Dillon; Barron, Jonathan T.. (2018). Unprocessing Images for Learned Raw Denoising. https://arxiv.org/abs/1811.11127.
14. Cameron, A. Colin, and Douglas L. Miller. “A Practitioner’s Guide to Cluster-Robust Inference.” *Journal of Human Resources*, vol. 50, no. 2, 2015, pp. 317-372. https://doi.org/10.3368/jhr.50.2.317.
15. Shannon, C.E.. (1949). Communication in the Presence of Noise. Proceedings of the IRE. https://doi.org/10.1109/jrproc.1949.232969.
16. Yang, Sidi, et al. "MANIQA: Multi-dimension Attention Network for No-Reference Image Quality Assessment." *arXiv preprint arXiv:2204.08958*, 2022. https://arxiv.org/abs/2204.08958.

17. Higham, Nicholas J.. (2002). Accuracy and Stability of Numerical Algorithms: Second Edition. Society for Industrial and Applied Mathematics. https://doi.org/10.1137/1.9780898718027.

18. Cook, R. Dennis. (1977). Detection of Influential Observation in Linear Regression. Technometrics. https://doi.org/10.1080/00401706.1977.10489493.

19. Harris, Charles R.; Millman, K. Jarrod; van der Walt, Stéfan J.; Gommers, Ralf; Virtanen, Pauli; Cournapeau, David; Wieser, Eric; Taylor, Julian; Berg, Sebastian; Smith, Nathaniel J.; Kern, Robert; Picus, Matti; Hoyer, Stephan; van Kerkwijk, Marten H.; Brett, Matthew; Haldane, Allan; del Río, Jaime Fernández; Wiebe, Mark; Peterson, Pearu; Gérard-Marchant, Pierre; Sheppard, Kevin; Reddy, Tyler; Weckesser, Warren; Abbasi, Hameer; Gohlke, Christoph; Oliphant, Travis E.. (2020). Array programming with NumPy. Nature. https://doi.org/10.1038/s41586-020-2649-2.

20. Meng, Zijie; Zeng, Yuanze; Chang, Xiang; Xu, Tianshuo; Chao, Fei; Cao, Xixin; Shang, Changjing; Shen, Qiang. (2025). Orpaint: a zero-shot inpainting model for oracle bone inscription rubbings with visual mamba block. Science China Information Sciences, 68(8), 189102. https://doi.org/10.1007/s11432-024-4493-4.

21. Wilkinson, Mark D.; Dumontier, Michel; Aalbersberg, IJsbrand Jan; Appleton, Gabrielle; Axton, Myles; Baak, Arie; Blomberg, Niklas; Boiten, Jan-Willem; da Silva Santos, Luiz Bonino; Bourne, Philip E.; Bouwman, Jildau; Brookes, Anthony J.; Clark, Tim; Crosas, Mercè; Dillo, Ingrid; Dumon, Olivier; Edmunds, Scott; Evelo, Chris T.; Finkers, Richard; Gonzalez-Beltran, Alejandra; Gray, Alasdair J.G.; Groth, Paul; Goble, Carole; Grethe, Jeffrey S.; Heringa, Jaap; ’t Hoen, Peter A.C; Hooft, Rob; Kuhn, Tobias; Kok, Ruben; Kok, Joost; Lusher, Scott J.; Martone, Maryann E.; Mons, Albert; Packer, Abel L.; Persson, Bengt; Rocca-Serra, Philippe; Roos, Marco; van Schaik, Rene; Sansone, Susanna-Assunta; Schultes, Erik; Sengstag, Thierry; Slater, Ted; Strawn, George; Swertz, Morris A.; Thompson, Mark; van der Lei, Johan; van Mulligen, Erik; Velterop, Jan; Waagmeester, Andra; Wittenburg, Peter; Wolstencroft, Katherine; Zhao, Jun; Mons, Barend. (2016). The FAIR Guiding Principles for scientific data management and stewardship. Scientific Data. https://doi.org/10.1038/sdata.2016.18.

22. Wang, Zhou, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. "Image Quality Assessment: From Error Visibility to Structural Similarity." *IEEE Transactions on Image Processing*, vol. 13, no. 4, 2004, pp. 600-612. https://doi.org/10.1109/TIP.2003.819861.

23. Mittal, A.; Moorthy, A. K.; Bovik, A. C.. (2012). No-Reference Image Quality Assessment in the Spatial Domain. IEEE Transactions on Image Processing. https://doi.org/10.1109/tip.2012.2214050.

24. Ding, Keyan, et al. "Image Quality Assessment: Unifying Structure and Texture Similarity." *arXiv preprint arXiv:2004.07728*, 2020. https://arxiv.org/abs/2004.07728.

25. Zhang, Richard, et al. "The Unreasonable Effectiveness of Deep Features as a Perceptual Metric" *arXiv preprint arXiv:1801.03924*, 2018. https://arxiv.org/abs/1801.03924.

26. Pratikakis, Ioannis; Zagori, Konstantinos; Kaddas, Panagiotis; Gatos, Basilis. (2018). ICFHR 2018 Competition on Handwritten Document Image Binarization (H-DIBCO 2018). 2018 16th International Conference on Frontiers in Handwriting Recognition (ICFHR). https://doi.org/10.1109/icfhr-2018.2018.00091.

27. Otsu, Nobuyuki. (1979). A Threshold Selection Method from Gray-Level Histograms. IEEE Transactions on Systems, Man, and Cybernetics. https://doi.org/10.1109/tsmc.1979.4310076.

28. Chen, Chen; Chen, Qifeng; Xu, Jia; Koltun, Vladlen. (2018). Learning to See in the Dark. https://arxiv.org/abs/1805.01934.

29. Virtanen, Pauli; Gommers, Ralf; Oliphant, Travis E.; Haberland, Matt; Reddy, Tyler; Cournapeau, David; Burovski, Evgeni; Peterson, Pearu; Weckesser, Warren; Bright, Jonathan; van der Walt, Stéfan J.; Brett, Matthew; Wilson, Joshua; Millman, K. Jarrod; Mayorov, Nikolay; Nelson, Andrew R. J.; Jones, Eric; Kern, Robert; Larson, Eric; Carey, C J; Polat, İlhan; Feng, Yu; Moore, Eric W.; VanderPlas, Jake; Laxalde, Denis; Perktold, Josef; Cimrman, Robert; Henriksen, Ian; Quintero, E. A.; Harris, Charles R.; Archibald, Anne M.; Ribeiro, Antônio H.; Pedregosa, Fabian; van Mulbregt, Paul; SciPy 1.0 Contributors; Vijaykumar, Aditya; Bardelli, Alessandro Pietro; Rothberg, Alex; Hilboll, Andreas; Kloeckner, Andreas; Scopatz, Anthony; Lee, Antony; Rokem, Ariel; Woods, C. Nathan; Fulton, Chad; Masson, Charles; Häggström, Christian; Fitzgerald, Clark; Nicholson, David A.; Hagen, David R.; Pasechnik, Dmitrii V.; Olivetti, Emanuele; Martin, Eric; Wieser, Eric; Silva, Fabrice; Lenders, Felix; Wilhelm, Florian; Young, G.; Price, Gavin A.; Ingold, Gert-Ludwig; Allen, Gregory E.; Lee, Gregory R.; Audren, Hervé; Probst, Irvin; Dietrich, Jörg P.; Silterra, Jacob; Webber, James T; Slavič, Janko; Nothman, Joel; Buchner, Johannes; Kulick, Johannes; Schönberger, Johannes L.; de Miranda Cardoso, José Vinícius; Reimer, Joscha; Harrington, Joseph; Rodríguez, Juan Luis Cano; Nunez-Iglesias, Juan; Kuczynski, Justin; Tritz, Kevin; Thoma, Martin; Newville, Matthew; Kümmerer, Matthias; Bolingbroke, Maximilian; Tartre, Michael; Pak, Mikhail; Smith, Nathaniel J.; Nowaczyk, Nikolai; Shebanov, Nikolay; Pavlyk, Oleksandr; Brodtkorb, Per A.; Lee, Perry; McGibbon, Robert T.; Feldbauer, Roman; Lewis, Sam; Tygier, Sam; Sievert, Scott; Vigna, Sebastiano; Peterson, Stefan; More, Surhud; Pudlik, Tadeusz; Oshima, Takuya; Pingel, Thomas J.; Robitaille, Thomas P.; Spura, Thomas; Jones, Thouis R.; Cera, Tim; Leslie, Tim; Zito, Tiziano; Krauss, Tom; Upadhyay, Utkarsh; Halchenko, Yaroslav O.; Vázquez-Baeza, Yoshiki. (2020). SciPy 1.0: fundamental algorithms for scientific computing in Python. Nature Methods. https://doi.org/10.1038/s41592-019-0686-2.

30. Candes, Emmanuel J; Plan, Yaniv. (2010). Matrix Completion With Noise. Proceedings of the IEEE. https://doi.org/10.1109/jproc.2009.2035722.