

RESEARCH

Open Access



Super-resolving star clusters with sheaves

Michael Robinson^{1,2*}  and Christopher T. Capraro²

*Correspondence:

michaelr@american.edu

¹ Department

of Mathematics and Statistics,

American University,

Washington, DC, USA

Full list of author information

is available at the end of the

article

Abstract

This article presents an optimization-based approach for counting and localizing stars within a small cluster, based on photon counts in a focal plane array. The array need not be arranged in any particular way, and relatively small numbers of photons are required to ensure convergence. The stars can be located close to one another, as the location and magnitude errors were found to be low when the separation was larger than 0.2 Rayleigh radii. To ensure generality of our approach, it was constructed as a special case of a general theory built upon topological signal processing using the mathematics of sheaves that is applicable to general source decomposition problems.

Keywords: Star cluster, Super-resolution, Sheaf, Topological signal processing

Mathematics Subject Classification: Primary 85-08

1 Introduction

This article presents an optimization-based approach for counting and localizing stars within a small cluster, based on photon counts in a focal plane array. The array need not be arranged in any particular way, and relatively small numbers of photons are required to ensure convergence. The stars can be located close to one another, and good performance was obtained when the separation was larger than 0.2 Rayleigh radii. To ensure generality of our approach, it was constructed as a special case of a general theory built upon topological signal processing using the mathematics of *sheaves*. This generality ensures that the strategy works without change for super-resolution problems on arbitrary topological spaces, including abstract simplicial complexes and other higher-order networks. We obtain theoretical performance guarantees for our method, in contrast to many state-of-the-art algorithms, because there are ultimately few parameters that the user of our method is required to choose in order to ensure good performance. While familiarity with the basics of sheaf-based topological signal processing is necessary to understand the derivation of the algorithm, this familiarity is not essential to understand the final form of algorithm.

1.1 Historical context

Perhaps the most famous algorithm for separating nearly coincident stars is the CLEAN algorithm and its generalizations (see [1–3], for instance, among many others). These greedy algorithms rely on the geometric structure of stars, namely that they are point

sources or nearly so. It should be noted that the more modern versions of CLEAN tend to require the user to select or “tune” parameters or function dictionaries to exploit subtle features of the scene, whereas the original CLEAN algorithm does not require this tuning since it assumes the stars are all point sources. In low lighting conditions, it is helpful to use statistical methods. For instance, Richardson–Lucy deconvolution can be effective at identifying star locations in the presence of shot noise [4, 5].

While greedy algorithms can yield good performance, they also can fail in dramatic ways. This can be especially pronounced if photon counts are low. This article proposes that an optimization-based approach can lend some robustness to an algorithm. Optimization-based imaging using the Wasserstein metric has recently led to improvements in molecular microscopy [6]. Our approach is a generalization of this idea.

This article uses a diagrammatic approach to topological signal processing, in which measurements of a signal are localized to portions of an arbitrary topological space using the formalism of a sheaf [7]. Due to their expressivity, sheaves have become popular tools for studying sensor and communication problems [8–12]. Our algorithm is an incremental minimization of sheaf *local consistency radius*, which is a general framework for solving sensing problems [13–15]. The key benefit of this framework is that once a parametric model of the signal has been constructed, parameter estimation from measurements is given by a canonical minimization problem. Although we apply it to a specific case, the methodology behind our algorithm is quite general. In a companion paper [16], we prove that our mathematical approach works for arbitrary sparse source decomposition problems.

1.2 Problem statement

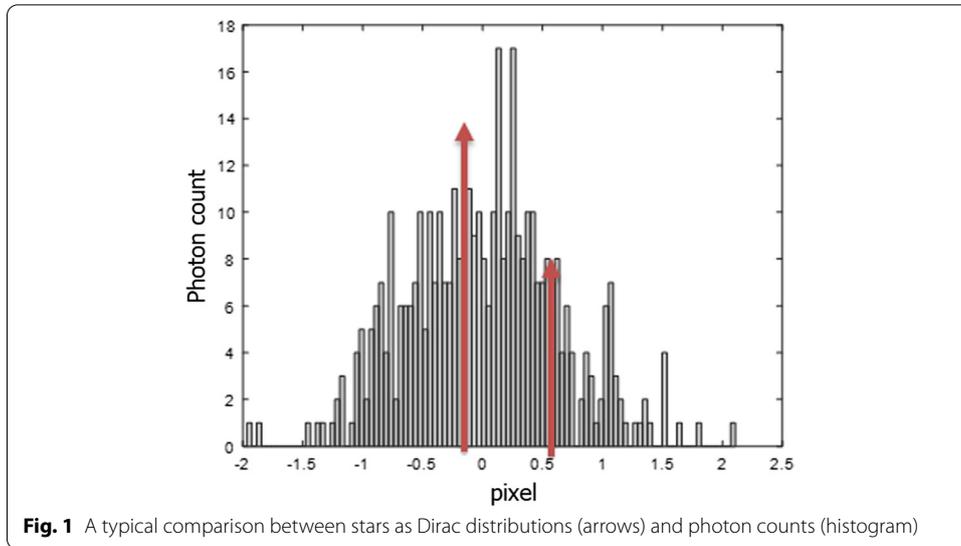
We can model a scene of N stars illuminating a focal plane by using a simplistic Dirac point-mass model,

$$s(x; a_1, b_1, a_2, b_2, \dots, a_N, b_N) = \sum_{n=1}^N a_n \delta(|x - b_n|), \quad (1)$$

where $a_n \in [0, \infty)$ is the magnitude of each star, $b_n \in \mathbb{R}^2$ is the location of each star in the focal plane, and $x \in \mathbb{R}^2$ can be taken to be the location of an arbitrary pixel. Although it is natural to assume that the pixels are organized in a grid, there is no actual requirement that this be so. Nevertheless, this article will assume a rectangular grid of pixels for simplicity of presentation.

The response of an actual sensor to the incoming photons from a set of stars in the scene will not be a sum of Dirac point-masses because of dispersion, diffraction, and other physical effects. Suppose that there are M pixels in the focal plane, each collecting an integer number of photons. This means that the measurements lie in \mathbb{N}^M (rather than \mathbb{R}^M). It is useful to record the photon counts as an integer-valued function $z = z(x) \in \ell^2(\mathbb{R}^2)$ in the focal plane \mathbb{R}^2 , as is shown schematically in Fig. 1. The problem to be solved in this article is therefore the determination of the number of stars N , the set $\{a_n\}$ (magnitudes), and the set $\{b_n\}$ (locations) from z (photon counts).

It is easy to imagine other signal decomposition problems that are modeled by Eq. (1). Instead of a superposition of stars, instead one can model the signal as a superposition



of parametric sources. All it takes is a different interpretation of the space of possible values for the locations b_n , and the replacement of the δ distribution with a prototypical signal from an appropriate dictionary. The mathematical implications—but not the practical utility—of this generalized approach are explored in [16].

2 Methods

Consider that each star is determined by its magnitude (an element of $[0, \infty)$) and location (an element of \mathbb{R}^2). These properties have different units and thereby suggest that the correct space for the parameters defining a star is

$$I = [0, \infty) \times \mathbb{R}^2 \subset \mathbb{R} \times \mathbb{R}^2.$$

There is a natural metric on I that is induced by the usual metric on \mathbb{R}^3 . Ideally, zero-magnitude stars should be indistinguishable from each other even if their locations differ. Although the induced metric does not satisfy this property, it will suffice for our purposes.

If there are N stars present in the scene, this would suggest that I^N is the correct space of parameters defining all stars. However, since the order in which the stars are listed in such a product does not matter, the correct space of parameters for defining all of the stars is the quotient space I^N/S_N , where S_N is the group of permutations on N items. Observe in particular that while I^N/S_N has the structure of a metric space, it is not a vector space!

2.1 Warm up: sheaf model for a known star count

Briefly, a *sheaf diagram* is a directed acyclic graph in which vertices are labeled with spaces of parameters (called *stalks*), and edges are labeled with functions (called *restrictions*) from the stalk shown at the tail of the edge to the stalk shown at its head. Sheaf diagrams have the additional requirement that the composition of consecutive functions along a path is independent of the path itself.

In this article, we will think of the sheaf diagram not as a graph, but as the Hasse diagram for a partially ordered set. Each vertex corresponds to an element of the partially ordered set, and each edge $a \rightarrow b$ corresponds to the relation $a \leq b$. This interpretation is helpful because every partially ordered set has a natural *Alexandrov topology*—a kind of higher-order network—in which each open set is the union of *upwardly closed* sets of the form

$$U_x = \{y : x \leq y\}$$

for each vertex x in the directed acyclic graph.

The specification of values from the stalks within an open set U is called an *assignment to U* for the sheaf. Assignments of values from every stalk that are *consistent* with the sheaf are called *global sections*; this means that $f(a) = b$ whenever $f : A \rightarrow B$ is a restriction in the sheaf diagram and $a \in A, b \in B$ are specified by the assignment. Intuitively, if a sheaf diagram represents a signal model, then the global sections characterize the signals permitted by that model.

If we know the number of stars is N , we can express the relationship between the photon count at each pixel x_k and the star magnitudes and locations by the sheaf diagram

$$\begin{array}{ccccc}
 & & I^N/S_N & & \\
 & \swarrow^{s(x_1;\cdot)} & \downarrow^{s(x_2;\cdot)} & \searrow^{s(x_M;\cdot)} & \\
 \mathbb{N} & & \mathbb{N} & \dots & \mathbb{N}
 \end{array} \tag{2}$$

The diagram expresses the fact that each photon count (values in the second row of the diagram) is determined by Eq. (1). Furthermore, no photon count functionally determines any other because there are no arrows between vertices in the second row. One should be aware that the lack of functional dependence does not imply that the photon counts are independent. Specifically, it may happen that one or more photon counts suffice to determine a unique value in I^N/S_N , after which point all of the other photon counts can be determined from this unique value. This is exactly the situation in which an adaptive method is warranted, since later photon counts are not actually required!

Suppose we have a set of M pixels at x_1, x_2, \dots, x_M , for which we have photon counts z_1, z_2, \dots, z_M respectively. These measurements form an assignment to the above sheaf, which is supported on the second row.

The assignment is not supported at the top row, since that would be tantamount to knowing all of the star parameters. Nevertheless, assuming that no noise is present and that we have correctly found the parameters $a_1, b_1, \dots, a_N, b_N$, then it should be the case that

$$z_m = s(x_m; a_1, b_1, \dots, a_N, b_N) \tag{3}$$

for all $m = 1, \dots, M$.

If M real measurements are taken from the scene determined by N stars, then the previous discussions mean that we can instead use a simpler sheaf diagram

$$\begin{array}{c} I^N/S_N \\ \downarrow f \\ \mathbb{N}^M \end{array}$$

where

$$f_N(a_1, b_1, \dots, a_N, b_N) := \begin{pmatrix} s(x_1; a_1, b_1, \dots, a_N, b_N) \\ s(x_2; a_1, b_1, \dots, a_N, b_N) \\ \vdots \\ s(x_M; a_1, b_1, \dots, a_N, b_N) \end{pmatrix}, \tag{4}$$

simply aggregates all of the pixel values into a single vector. One can consider this as a function taking specific values on regions within the plane \mathbb{R}^2 , so we can interpret $f_N : I^N/S_N \rightarrow \ell^2(\mathbb{R}^2)$.

In the (unlikely) situation that the pixels have taken exactly the values specified by Eq. (1), this corresponds to a global section of both of the sheaf diagrams shown above. The presence of noise or systematic errors implies that Eq. (3) will generally not hold exactly. The sheaf *consistency radius* measures the degree to which Eq. (3) holds and can be computed for any sheaf in which the stalks are metric spaces [15]. The consistency radius in our case is merely¹ a rearrangement of Eq. (3), namely

$$c(a_1, b_1, \dots, a_n, b_n) = \sqrt{\sum_{m=1}^M |s(x_m; a_1, b_1, \dots, a_n, b_n) - z_m|^2}.$$

Since global sections have zero consistency radius and the stalks of the sheaf are all metric spaces, any assignment which is not a global section will have positive consistency radius.

On the other hand, if the f_N function defined in Eq. (4) is injective then there is only one global section. This means that we can recast the problem of obtaining the stars from the photon counts as the minimization of consistency radius. With essentially no further work, we obtain the following result.

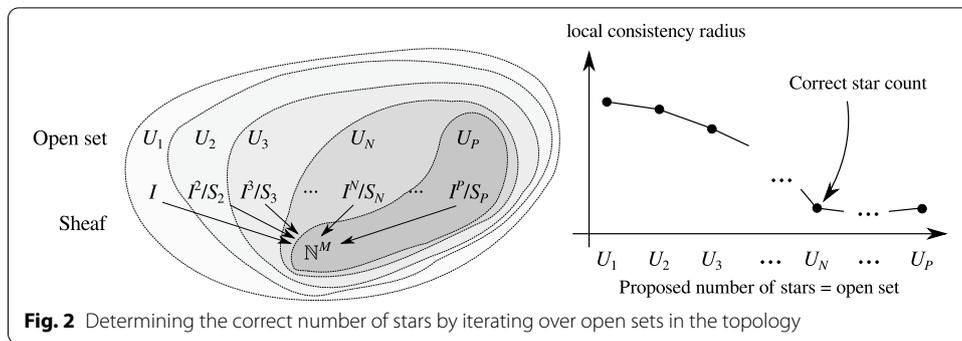
Corollary 1 *If the function f_N defined in Eq. (4) is injective, then the solution to*

$$\operatorname{argmin}_{\{a_n, b_n\}_{n=1}^N \in I^N/S_N} \left(\sum_{m=1}^M |s(x_m; a_1, b_1, \dots, a_n, b_n) - z_m|^2 \right). \tag{5}$$

is the correct star decomposition for the scene $\{z_m\}_{m=1}^M$ if one exists.

The reader may correctly argue that the optimization problem in Eq. (5) can be obtained easily—though not solved—by inspection! However, if the number of stars is not known, then the correct optimization problem that obtains the star parameters is

¹ The reader is cautioned that the consistency radius for more complicated sheaves can be difficult to determine by inspection!



difficult to state explicitly. However, the correct optimization problem is still a consistency radius minimization, but for a more elaborate sheaf. This sheaf accounts both for the unknown number of stars and the fact that the number of measurements that should be taken depends on the (unknown) number of stars.

2.2 Sheaf model for an unknown star count

Now let us assume that the scene is built from an unknown number of stars. We cannot use a single f_N map, because we do not know N . Therefore, we need to consider many possibilities for the number of stars. For instance, let us consider that there may be as many as P stars in the scene.

Definition 1 The sheaf \mathcal{J}_P is defined by the diagram

$$\begin{array}{ccccccc}
 I & & I^2/S_2 & & I^3/S_3 & \dots & I^P/S_P \quad (6) \\
 & \searrow^{f_1(\cdot)} & \downarrow^{f_2(\cdot, \cdot)} & \searrow^{f_3(\cdot, \cdot, \cdot)} & & & \\
 & & \mathbb{N}^M & & & &
 \end{array}$$

in which each of the restriction maps is given by the map defined in Eq. (4) with different numbers of stars.

Notice that the P in \mathcal{J}_P is the maximum number of stars we will attempt to consider, which may not have much (if anything) to do with the actual number of stars N . Of course, we hope that $N \leq P$ so that the actual number of stars is considered in our analysis.

Since the underlying directed acyclic graph of \mathcal{J}_P is more complicated than in the previous section, it is useful to consider what the open sets are. While there are many open sets, for our purposes the ones we need to consider are unions of the U_k sets shown in the left frame of Fig. 2. The set U_k corresponds to a situation where we posit that there are between k and P stars in the scene, and the open set U_P corresponds to a case where we assume that there are exactly P stars in the scene.

An assignment in the case of U_P is given by an element of I^P/S_P paired with an element of \mathbb{N}^M , namely a set of the form $\{(a_1, b_1), \dots, (a_P, b_P)\}$ paired with a $z \in \mathbb{N}^M$. An assignment to an arbitrary U_k is a bit more complicated, since it contains an element z of

\mathbb{N}^M along with an element of each of $I^k/S_k, \dots, I^P/S_P$. To keep the indexing clear, let us write $w_{k,\bullet} \in I^k/S_k$ for each of these elements, realizing that $w_{k,j}$ is a magnitude/location pair for each $j = 1, \dots, k$.

Given this notation, the global consistency radius of an assignment to this sheaf is given by

$$c((w_{\bullet,\bullet}, z)) = \sum_{i=1}^P \|f_i(w_{i,1}, \dots, w_{i,i}) - z\|,$$

where $z \in \mathbb{N}^M$ is the set of photon counts and the $w_{i,j}$ is the proposed j -th star parameter when we are considering a representation of the scene with i stars. The global consistency radius is simply the sum of consistency radii of each subproblem, in which we consider a given number of stars.

On the other hand, each one of these subproblems is the minimization of the consistency radius restricted to each open set in the base space topology for \mathcal{J}_P . When restricted to an open set U_k , the resulting *local* consistency radius is

$$c((w_{\bullet,\bullet}, z); U_k) = \sum_{i=k}^P \|f_i(w_{i,1}, \dots, w_{i,i}) - z\|,$$

and cannot exceed the global consistency radius [17].

If the f_\bullet functions are injective, then it is clear that if $N > P$, the only set upon which the local consistency radius can vanish is the set consisting of only the photon counts (i.e., the bottom row of the sheaf diagram). Conversely, if the f_\bullet functions are injective and $N \leq P$, then the minimum local consistency radius possible on U_P is zero.

2.3 Performance guarantees

Proposition 1 *Suppose that the number of stars is fixed at N and that $z \in \mathbb{N}^M$ is given by*

$$z = S(a_1, b_1, \dots, a_N, b_N),$$

where S is given by Eq. (4).

If we assign $z \in \mathbb{N}^M$ to represent the set of photon counts in the sheaf \mathcal{J}_P (see Eq. 6), then there is an extension to a global assignment in which the local consistency radius vanishes for all sufficiently small open sets provided that $P \geq N$.

Proof To that end, $z \in \mathbb{N}^M$ is assigned in the bottom row of the diagram in Eq. (6).

Let us name the values in the assignment to the top row of the sheaf diagram

$$\{(a_{1,1}, b_{1,1})\}, \{(a_{2,1}, b_{2,1}), (a_{2,2}, b_{2,2})\}, \{(a_{3,1}, b_{3,1}), (a_{3,2}, b_{3,2}), (a_{3,3}, b_{3,3})\}, \dots, \{(a_{P,1}, b_{P,1}), \dots, (a_{P,P}, b_{P,P})\}.$$

Notice that these values define many possibilities for the source parameters.

If the number of stars is not known, but decompositions with $P \geq N$ stars are encoded in the sheaf, then the consistency radius of any assignment supported on the open set \mathcal{U}_N is given by

$$\sum_{i=N}^P \|S(a_{i,1}, b_{i,1} \dots, a_{i,i}, b_{i,i}) - z\|. \tag{7}$$

Define

$$a_{P,k} = \begin{cases} a_k & \text{if } k \leq N, \\ 0 & \text{otherwise} \end{cases}$$

and

$$b_{P,k} = \begin{cases} b_k & \text{if } k \leq N, \\ \text{arbitrary} & \text{otherwise.} \end{cases}$$

If we then declare that $a_{i,k} = a_{j,k}$ and $b_{i,k} = b_{j,k}$ whenever they are both defined, then this definition clearly results in the sum of Eq. (6) being zero. \square

Corollary 2 *If the S function is injective, then the assignment for which the local consistency radius of \mathcal{J}_P vanishes occurs precisely at the number of stars N .*

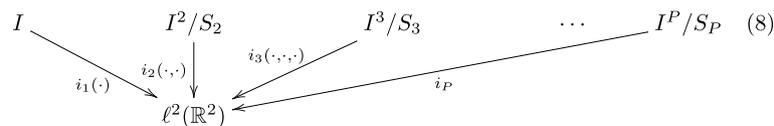
2.4 Probabilistic treatment

If we combine the sheaf models with the observation that the collection of photon counts from each pixel are best thought of as a (single) function on the plane, we obtain a way to recover star parameters from a distributional measurement.

To formalize this idea, consider the space $\ell^2(\mathbb{R}^2)$ of functions on star locations. In brief, a measurement specifies the amount of mass on a set of star locations. The map $i_N : I^N/S_N \rightarrow \ell^2(\mathbb{R}^2)$ reinterprets stars as measures

$$i_N((a_1, b_1), \dots, (a_N, b_N)) := \sum_{n=1}^N a_n \delta(x - b_n).$$

We can use these i_N maps as the restriction maps in \mathcal{J}_P , which still works structurally, namely



Algorithmically, if we obtain a measurement as an element $z \in \ell^2(\mathbb{R}^2)$ (which need not be a sum of a small number of Dirac distributions), then the best star decomposition is still obtained by minimizing local consistency radii:

$$\|i_P(w_1, \dots, w_P) - z\|^2,$$

and selecting the smallest P for which this quantity is zero. In cases where the source magnitudes are widely varying, it is likely that a greedy approach is possible (and may be preferable). The trade-off between greedy and non-greedy approaches to solving the minimization problem is quite subtle! Moreover, if one considers a generalized signal model, rather than a star cluster, the topology of the sheaf diagram can impact how this trade-off operates. The interested reader should consult [16] for details about the general case. In the next section, we employ a compromise approach by reusing solutions with smaller number of proposed stars when solving for larger numbers of proposed stars.

2.5 Pivoting to an algorithm

From a methodological perspective, Corollary 2 implies that the number of stars can be determined from the consistency filtration given enough samples. If there is noise, our algorithm looks for the “knee” in local consistency radius as the open set size increases, as shown in Fig. 2.

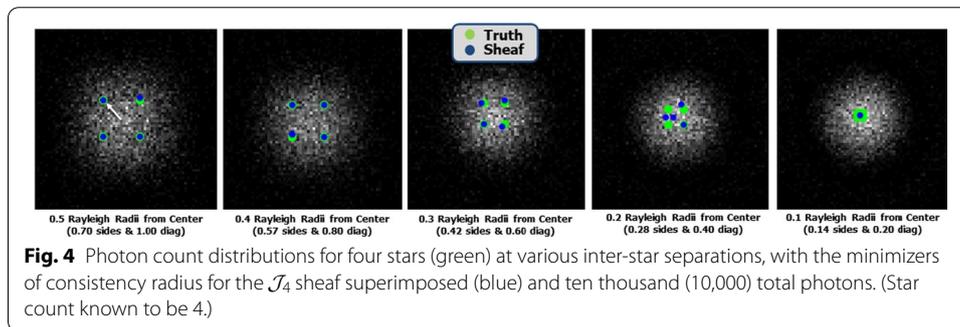
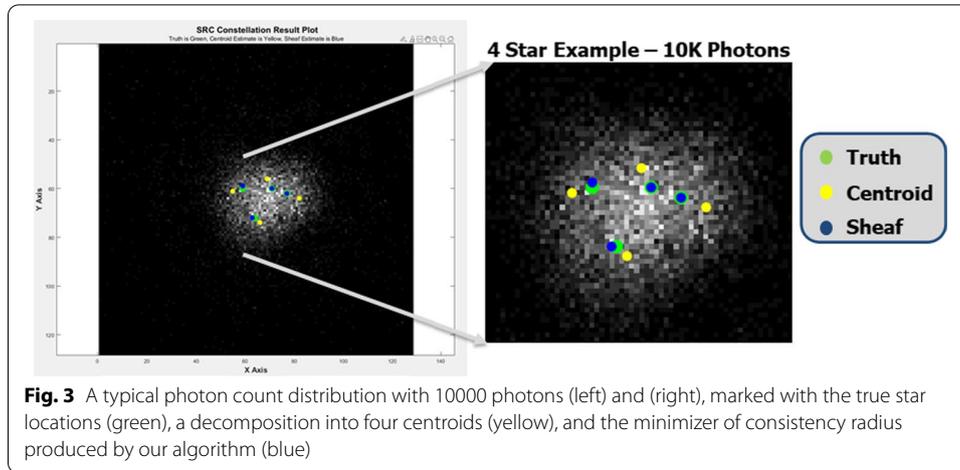
The algorithm we used to count, locate, and characterize stars begins with a proposed maximum number of stars P . Convergence is not ensured by Proposition 1 and Corollary 2 if the true number of stars N is larger than P . The algorithm proceeds through the following steps, which ultimately minimize the local consistency radii of the sheaf \mathcal{J}_P :

1. Collect photon counts from all pixels into a vector z
2. For i in $1, 2, \dots, P$,
 - (a) Initial guess: If $i = 1$, use the centroid of the photon count distribution as the initial guess for b_1 and the total photon count as a_1 . Otherwise, use the previous step’s set of star locations b_1, \dots, b_{i-1} and magnitudes a_1, \dots, a_{i-1} as an initial guess with $a_i = 0$ and b_i chosen randomly.
 - (b) Iteration: Solve the minimization problem Eq. (5) given i stars using the i centroids as initial guesses for b_1, \dots, b_i .
 - (c) Store: the residual error, along with the parameters for the i stars.
3. Return the proposed number of stars, locations, and magnitudes corresponding to the smallest residual.

We found that the minimization problem Eq. (5) was solvable using the standard Matlab `fmincon` function, though the initial guess stage was necessary to aid in ensuring convergence under repeated Monte Carlo runs.

3 Results and discussion

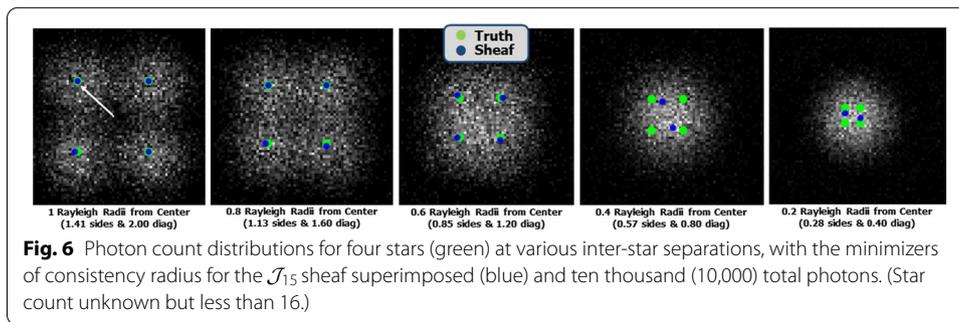
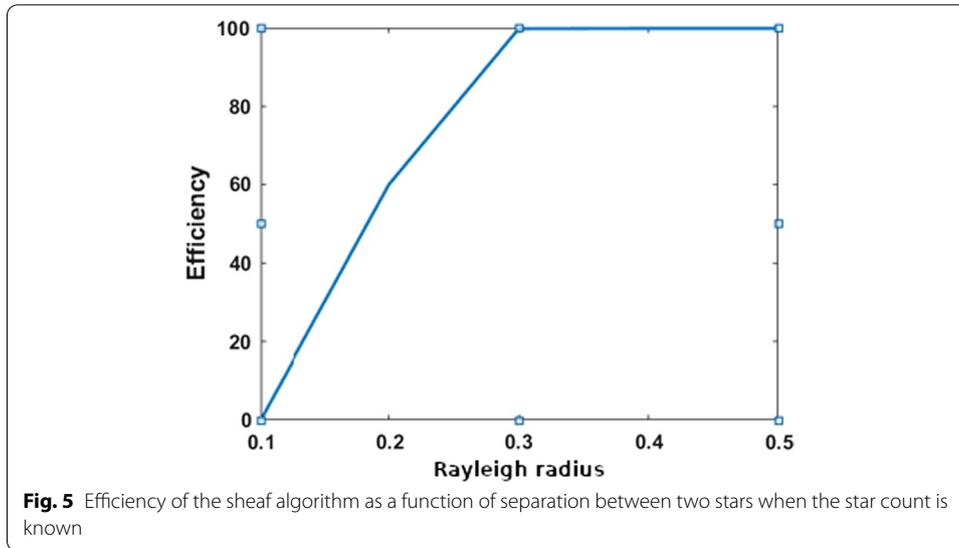
As a test of our approach, star fields consisting of between 1 and 10 stars were simulated. These star fields were propagated through a simulated optical system, resulting in photon counts on an array of pixels. These photon counts were then used to determine star locations and magnitudes using the algorithm discussed above, a multiple centroiding approach, the standard CLEAN algorithm [1], and Richardson–Lucy deconvolution [4, 5].



We chose to focus our attention on only these four algorithms, because they do not require the user to “tune” parameters in order to get good performance. In algorithms that require tuning, good performance is difficult to guarantee in both theory and practice, and this tuning may require considerable effort. In contrast, the four methods considered in our study do not require any such tuning, and can be compared directly with no further manipulation.

While the CLEAN algorithm, the multiple centroiding algorithm, and our sheaf algorithm all output star locations and brightness directly, the Richardson–Lucy algorithm produces in a denoised image. To convert this into star locations, we ran two iterations of the Richardson–Lucy algorithm, followed by a narrow Gaussian filters, a threshold detector, and then into the multiple centroid detector. This yielded a simple algorithm with only one tunable parameter (the Gaussian filter bandwidth), which we tuned for maximum *efficiency* as described below in Eq. (9).

Figure 3 shows a typical result for 10,000 photons. Figure 4 shows several example runs for four stars, in which the stars (shown in green) are placed on the corners of a square whose side length was varied. If we minimize consistency radius using \mathcal{J}_4 , which effectively means that we are declaring the number of stars to be known, the minimizers of consistency radius for \mathcal{J}_4 are shown in blue. Our algorithm does not expect any particular geometry, as is clear by its obvious symmetry breaking in the right frames of Fig. 4.



Assuming that the correct magnitude and location of star i are given by (a_i, b_i) and that the algorithm’s estimate of the same is (a'_i, b'_i) , an algorithm can be scored by its *efficiency*, given by

$$E := 100 - \min_{\phi} \sqrt{\left(1 - \frac{\#\text{image } \phi}{P}\right)^2 + \sum_{i=1}^N \sum_{j=1}^P |a_{\phi(j)} - a'_i|^2 + \sum_{i=1}^N \sum_{j=1}^P |b_{\phi(j)} - b'_i|^2}, \tag{9}$$

where $\phi : \{1, 2, \dots, P\} \rightarrow \{1, 2, \dots, N\}$ ranges over all functions. Note that $\#\text{image } \phi/P$ is the Jaccard index of the two sets of stars, given the function ϕ used as a matching. Higher efficiency means that the algorithm is doing a better job of determining the correct number of stars, their magnitudes, and locations. The best possible efficiency is 100, while the worst is 0.

Good efficiency scores are produced by our algorithm for separations greater than 0.2 Rayleigh radii when the star count is known, as is shown in Fig. 5. As might be expected, the performance of the algorithm suffers without knowing the correct number of stars. Figure 6 shows the same four stars, but with the minimizers of consistency radius for \mathcal{J}_{15} . This means that the algorithm was not told the correct number of stars, but merely that there are at most 15 stars. The overall efficiency scores are less in this case. As shown in Fig. 7, performance is good for separations greater than 0.5 Rayleigh radii.

Table 1 Constellation tests performed

Test	Star count	Magnitudes
A	1	N/A
B	2	Equal
C	2	Variable
D	3	Equal
E	3	Variable
F	7	Variable
G	10	Variable

shown in Fig. 9. For the other tests, the sheaf algorithm's performance is substantially more variable than CLEAN. Although the Richardson–Lucy algorithm does better than CLEAN in all cases, the sheaf does better than that by a similar margin. Interestingly, the pattern of which tests had higher efficiencies with the sheaf algorithm is similar to that of the Richardson–Lucy algorithm.

The location errors for all three algorithms and all tests are shown in Fig. 10. As should be expected, location errors increase in all three algorithms as the number of stars increase. Generally, the sheaf algorithm yields better location estimates than both CLEAN and the Richardson–Lucy for a smaller number of stars, while CLEAN yields better location estimates than the sheaf for a larger number of stars.

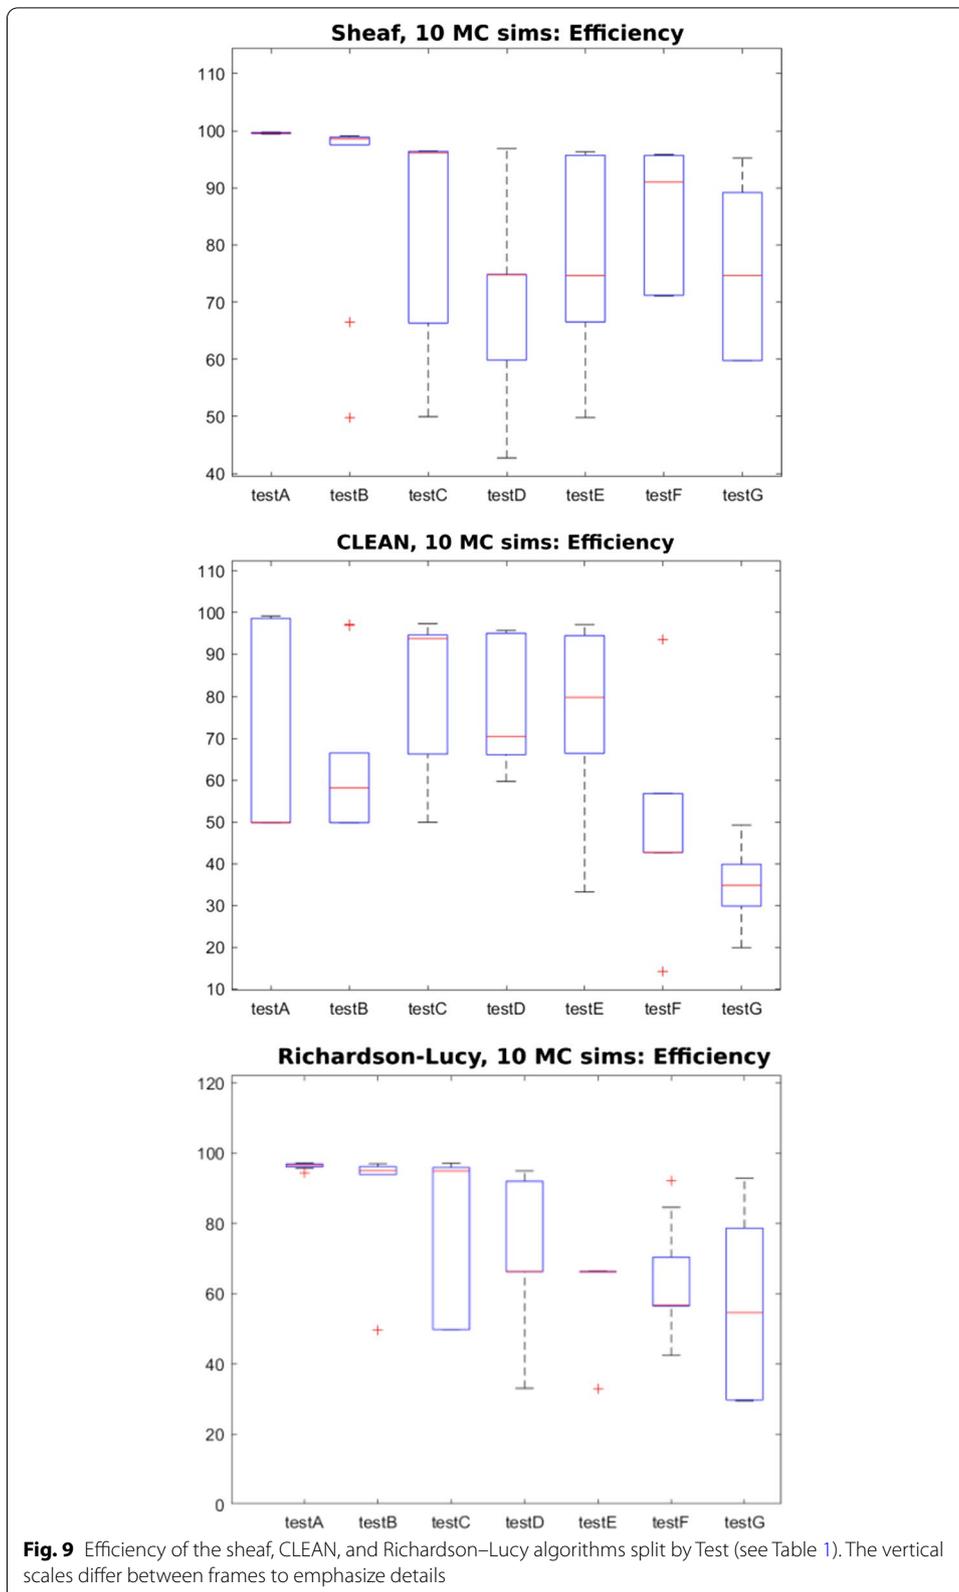
Magnitude estimate errors for all three algorithms and all tests are shown in Fig. 11. All three algorithms yield highly variable performance on magnitude estimates regardless of test conditions. At least in the sheaf algorithm case, this may be a direct result of the use of the Euclidean metric on $\ell^2(\mathbb{R}^2)$, since—in contrast to the Wasserstein metric [6]—the convergence in magnitude is not guaranteed.

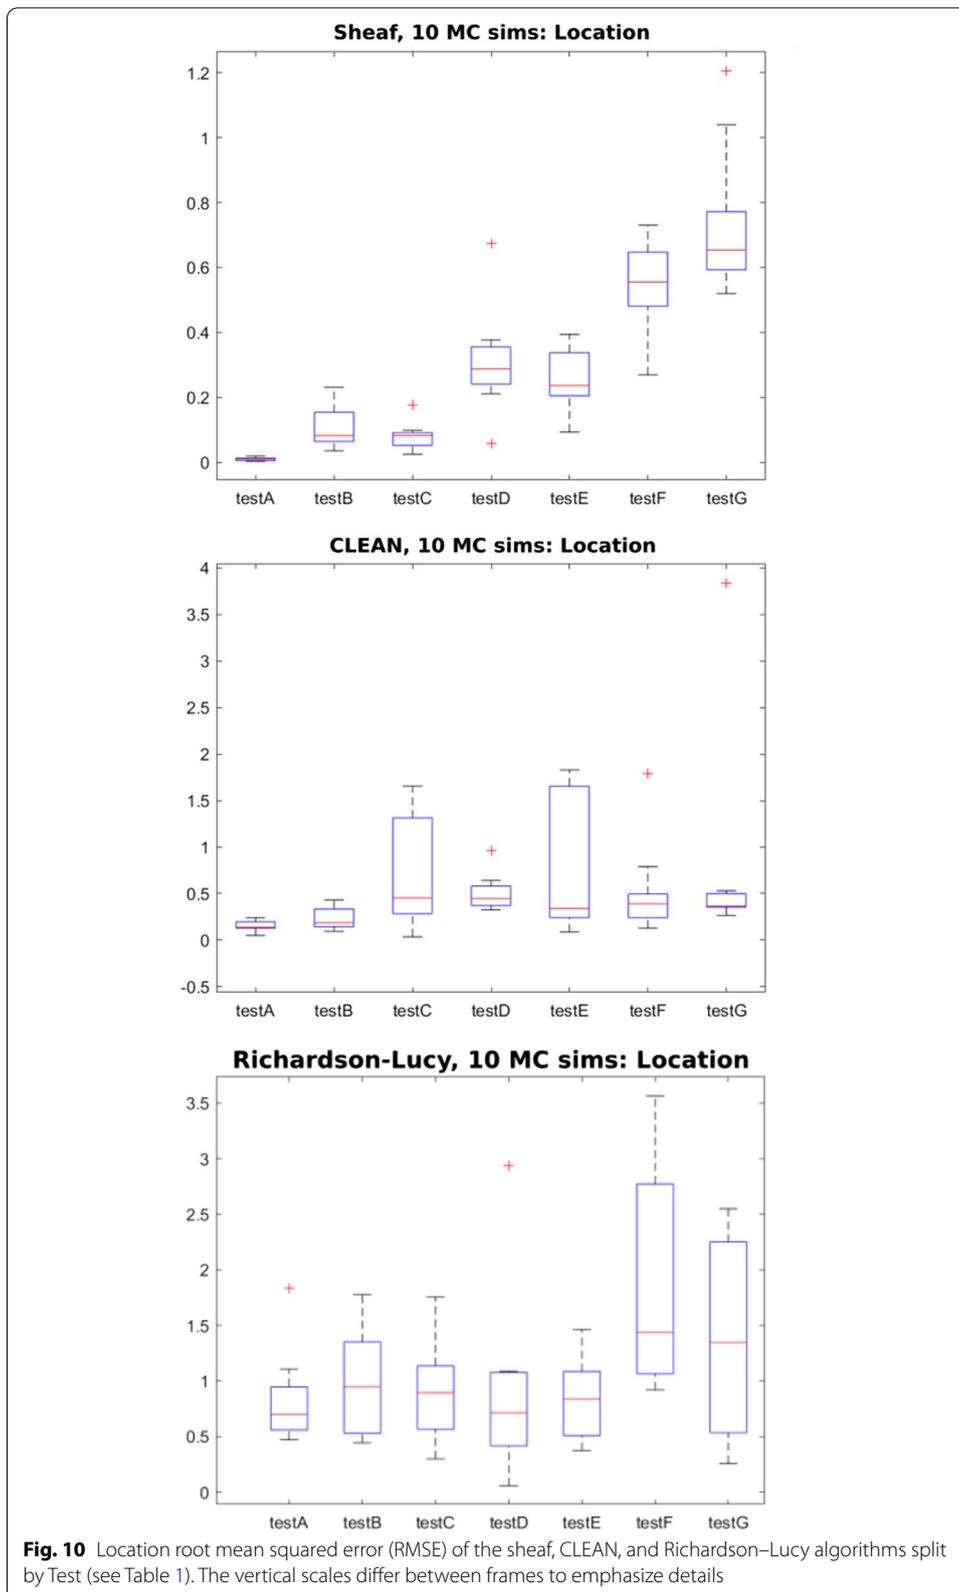
4 Conclusion

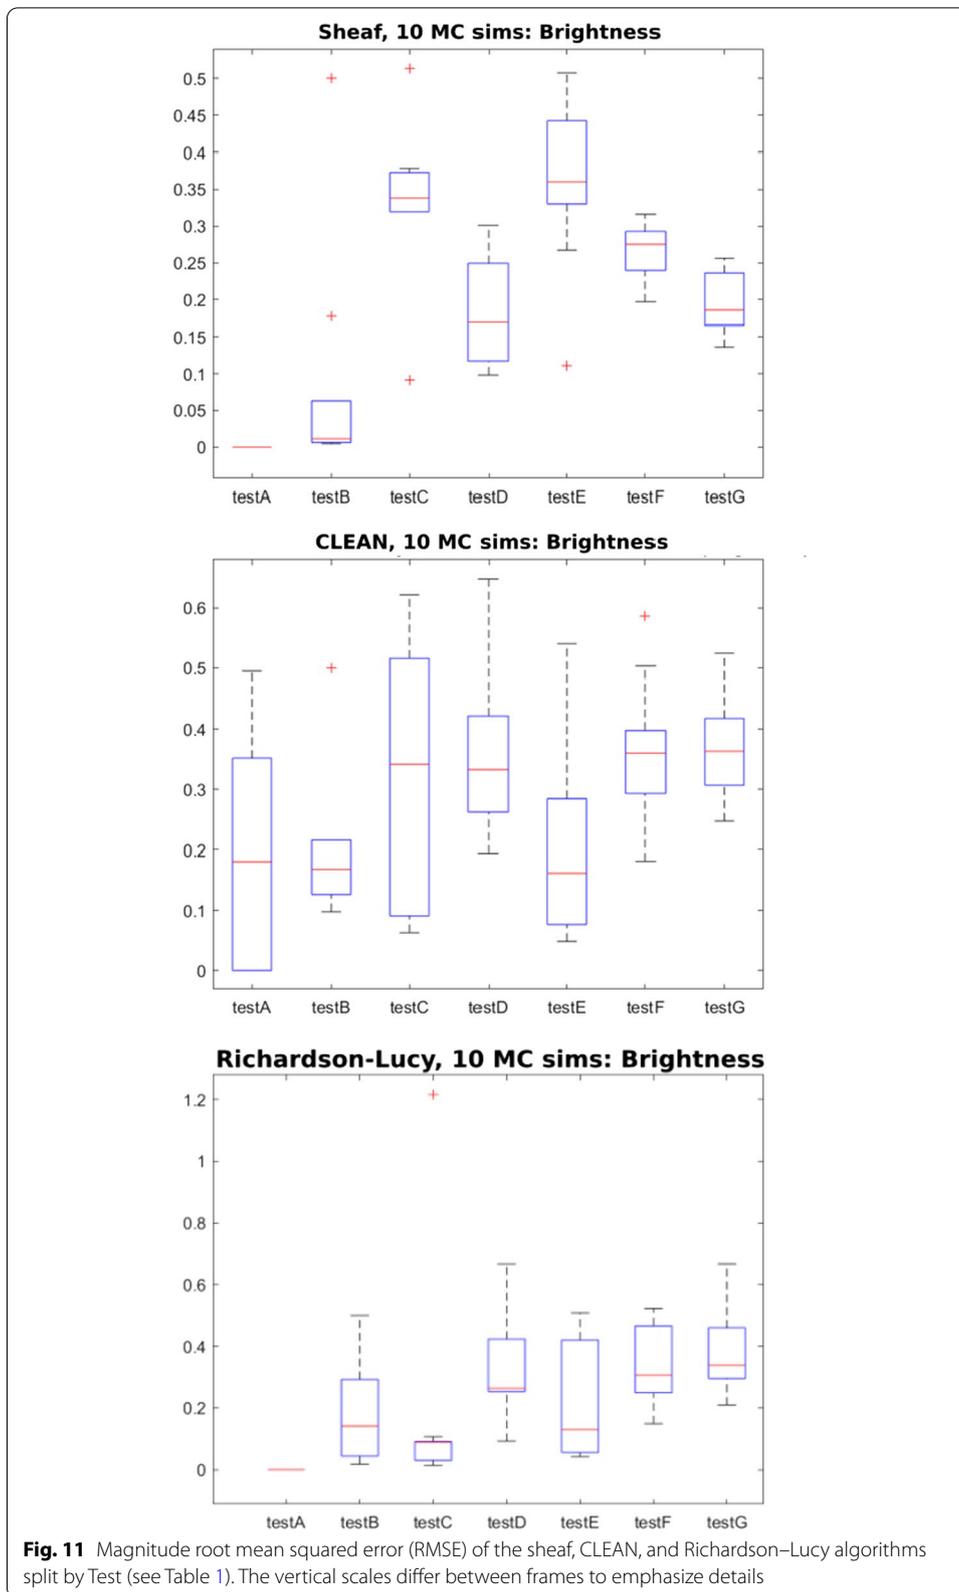
The results we obtained from our algorithm typically exceeded the performance of both the CLEAN and Richardson–Lucy algorithms, especially with more challenging scenes, which indicates that our sheaf method shows promise. Moreover, since our sheaf method does not require the user to tune parameters, it is easy to apply directly to a scene of interest. The fact that it converges very well for small numbers of stars and less well for larger numbers of stars suggests there is room for improvement. In particular, the use of a general optimization solver is not particularly efficient, and the solver converges slowly.

Computational considerations were not central to our approach on this problem, but could be important in future work. We have already developed a general library for sheaf algorithms [18] and could port this library to a high-performance computing platform. We expect that high-performance sheaf computational tools would have broad applicability, providing numerous avenues for improvements.

Finally, the sheaf-based approach we used is extremely general and applies without change to many other imaging techniques, not just focal plane arrays, as discussed in [16]. Quantum imaging with a small number of photons [19] fits neatly into this perspective and deserves to be explored in future work.







Acknowledgements

The authors thank Tom Ruekgauer and his team for the use of the simulated optical data presented in this article.

Authors' contributions

MR developed the theoretical framework and wrote the initial manuscript draft. CC developed the implementation of the algorithm and tested it against the simulated data. Both authors read and approved the final manuscript.

Funding

The views, opinions and/or findings expressed are those of the author and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government. This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Agreement No. HR00112090125.

Availability of data and materials

The data that support the findings of this study are available from Tom Ruekgauer, JHU/APL but restrictions apply to the availability of these data, which were used under license for the current study, and so are not publicly available. Data are, however, available from the authors upon reasonable request and with permission of Tom Ruekgauer.

Declarations

Ethics approval and consent to participate

Not applicable.

Consent for publication

Not applicable.

Competing interests

Michael Robinson and Chris Capraro declare that they have no competing interests.

Author details

¹Department of Mathematics and Statistics, American University, Washington, DC, USA. ²SRC, Inc., Syracuse, NY, USA.

Received: 1 December 2021 Accepted: 15 March 2022

Published online: 28 March 2022

References

1. J.A. Högbom, Aperture synthesis with a non-regular distribution of interferometer baselines. *Astron. Astrophys. Suppl.* **15**, 417 (1974)
2. G. McKinnon, C. Burger, P. Boesiger, Spectral baseline correction using CLEAN. *Mag. Resonan. Med.* **13**(1), 145–149 (1990)
3. I.M. Stewart, D.M. Fenech, T.W.B. Muxlow, A multiple-beam CLEAN for imaging intra-day variable radio sources. *Astron. Astrophys.* **535**, A81 (2011)
4. D. Whisler, Computational star-tracking compensation in astrophotography (2020)
5. L. Su, X. Shao, L. Wang, H. Wang, Y. Huang, Richardson–Lucy deblurring for the star scene under a thinning motion path, in *Commercial + Scientific Sensing and Imaging* (2015)
6. H. Mazidia, T. Dinga, A. Nehoraia, M.D. Lewa, Quantifying accuracy and heterogeneity in single-molecule super-resolution microscopy. *Nat. Commun.* **11**(6353), 1–11 (2020)
7. M. Robinson, *Topological Signal Processing* (Springer, Berlin, 2014)
8. R. Short, A. Hylton, R. Cardona, R. Green, G. Bainbridge, M. Moy, J. Cleveland, Towards sheaf theoretic analyses for delay tolerant networking, in *2021 IEEE Aerospace Conference (50100)* (2021), pp. 1–9
9. G. Essl, Topological IIR filters over simplicial topologies via sheaves. *IEEE Signal Process. Lett.* **27**, 1215–1219 (2020)
10. J. Hansen, R. Ghrist, Learning sheaf Laplacians from smooth signals, in *ICASSP 2019—2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (2019), pp. 5446–5450
11. M.A. Nguemo, C. Tcheka, Sheaf cohomology on network codings: maxflow-mincut theorem. *Appl. Gen. Topol.* **18**, 219–230 (2017)
12. R. Ghrist, S. Krishnan, A topological max-flow-min-cut theorem, in *2013 IEEE Global Conference on Signal and Information Processing* (2013), pp. 815–818
13. M. Robinson, Hunting for foxes with sheaves. *Not. Am. Math. Soc.* **66**(5), 661–676 (2019)
14. M. Robinson, J. Henrich, C. Capraro, P. Zulch, Dynamic sensor fusion using local topology, in *IEEE Aerospace Conference* (2018)
15. M. Robinson, Sheaves are the canonical data structure for sensor integration. *Inform. Fusion* **36**, 208–224 (2017)
16. M. Robinson, Aggregation sheaves for greedy modal decompositions (2021), arXiv:2106.04445
17. M. Robinson, Assignments to sheaves of pseudometric spaces. *Compositionality* **2**, 2 (2020)
18. M. Robinson, PySheaf: the Python sheaf library (2020). <https://github.com/kb1dds/pysheaf/>
19. M. Tsang, Resolving starlight: a quantum perspective. *Contemp. Phys.* **60**(4), 279–298 (2019)

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.