

Globally optimal and scalable N -way matching of astronomy catalogs

TU NGUYEN,¹ AMITABH BASU,^{1,2} AND TAMÁS BUDAVÁRI^{1,2,3}

¹*Dept. of Applied Mathematics & Statistics, Johns Hopkins University, Baltimore MD 21218*

²*Dept. of Computer Science, Johns Hopkins University, Baltimore MD 21218*

³*Dept. of Physics & Astronomy, Johns Hopkins University, Baltimore MD 21218*

ABSTRACT

Building on previous Bayesian approaches, we introduce a novel formulation of probabilistic cross-identification, where detections are directly associated to (hypothesized) astronomical objects in a globally optimal way. We show that this new method scales better for processing multiple catalogs than enumerating all possible candidates, especially in the limit of crowded fields, which is the most challenging observational regime for new-generation astronomy experiments such as the Rubin Observatory Legacy Survey of Space and Time (LSST). Here we study simulated catalogs where the ground-truth is known and report on the statistical and computational performance of the method. The paper is accompanied by a public software tool to perform globally optimal catalog matching based on directional data.

1. MOTIVATION

Several approaches have been proposed over the years to combine observations across telescopes and epochs. In the era of time-domain astronomy where dozens or hundreds of observation epochs are available, these problems are more important than ever. In particular, combining catalogs has been a central issue where detections in separate exposures are matched by identifying the ones that correspond to the same celestial object. Several tools were developed to provide solution to the cross-matching problem, such as TOPCAT (Taylor 2005) and CDS XMatch (Pineau et al. 2011; Boch et al. 2012). However, they do not consider the statistical aspect of the problem.

This cross-identification problem was successfully addressed using Bayesian hypothesis testing by Budavári & Szalay (2008), whose methodology was implemented in the latest version of the SkyQuery service (Budavári et al. 2013) which is now part of the SciServer Science Platform (Taghizadeh-Popp et al. 2020). The Bayesian formalism and the combinatorial nature of the problem is discussed in a review by Budavári & Loredo (2015). The first solution came from Budavári & Basu (2016) who formulated the matching problem as a search for globally optimal associations using combinatorial opti-

mization, where the marginal likelihood of the entire matched catalog is maximized, and used the Hungarian algorithm (Munkres 1957) to solve it. After that proof of concept was developed for two catalogs, Shi et al. (2019) extended the algorithm to handle multiple catalogs using Integer Linear Programming, or ILP for short. For simplicity, we will refer this method as *CanILP*, as it enumerates all possible candidate associations and uses ILP to choose the best valid subset. As we will discuss later, the method suggested in Shi et al. (2019) does not scale very well with large number of catalogs. This scaling problem is also observed in Pineau et al. (2017) as the authors try to estimate the probability, for all combinations of sources, that a tuple of sources from different catalogs correspond to the same object. The exhaustive search results in an exponential growth in the number of possible tuples as the number of catalogs increases. They also note that this approach is not feasible in practice for more than 9 catalogs.

In this paper, we improve on the previous studies by introducing a novel formulation, hereafter referred to as *DirILP*, where we use ILP to directly assign detections to hypothesized objects. Section 2 describes the new approach, and Section 3 illustrates how the new method scales better with the number of input catalogs. Section 4 discusses a public software tool to solve the catalog matching problem. Section 5 concludes the study.

2. OUR APPROACH

To quantify the associations among independent detections, a relatively recent approach was developed that uses a hierarchical Bayesian formalism. Suppose there are C catalogs, indexed by $c \in \{1, \dots, C\}$, with each catalog capturing N_c sources respectively. Let D_{ic} denote the measurements for source i in catalog c , hereafter denoted by tuple (i, c) . Associated with any (i, c) measurement is a likelihood function $\ell_{ic}(\omega) = p(D_{ic}|\omega)$, for the unknown true direction ω , which captures the astrometric uncertainty. While other object properties could also be considered in general, such as their brightness or colors, here we focus on directional matching only.

Here we adopt the definition for matching used in previous papers (e.g., Budavári & Szalay 2008; Budavári & Loredó 2015; Budavári & Basu 2016; Shi et al. 2019) where one tests whether two or more detections correspond to the same physical object. That said, it is possible to define the “match” hypothesis such that detections are not of the “same” object but instead just “part of” another, e.g., an optical galaxy in an X-ray cluster, or a star in a blend of two. In theory such scenarios can be accommodated (by changing the marginal likelihood calculations, see below), but the computation requirements might increase, and the interpretation of the matched catalog would be more difficult.¹ Our association approach described below is flexible and will yield matches based on the underlying model.

Associations are created by grouping all sources in all catalogs such that each belongs to only one group. Mathematically, a *partition* P is created of the data set D , the union of all sources in all catalogs, where each subset corresponds to the same celestial object. The number of subsets in the partition will constitute the number of hypothesized objects N_{obj} , which is an unknown but bounded quantity that is less or equal to the number of all sources. Typically it is much less as the equality would mean that every source is in fact a separate object altogether. We can index every object by an integer $o \in \{1, \dots, N_{\text{obj}}\}$. Let S_o be the set of sources (i, c) associated with object o and C_o be the list of catalogs containing sources associated with object o . Following Budavári & Loredó (2015), the likelihood of a partition P of all sources, a collection of the S_o subsets, will be a product of conditionally independent terms,

$$\mathcal{L}(P) \equiv p(D|P) = \prod_o \mathcal{M}_o, \quad (1)$$

¹ Further complications emerge when different blends are to be matched, in which case one considers whether the detections “share” components, e.g., a common star in two different blends.

where the marginal likelihood \mathcal{M}_o for the association corresponding to object o is

$$\mathcal{M}_o = \int d\omega \rho_{C_o}(\omega) \prod_{(i,c) \in S_o} \ell_{ic}(\omega). \quad (2)$$

Here $\rho_{C_o}(\omega)$ is the prior probability density function of the object direction producing sources within (the footprint of) every catalog in the set C_o . This notation enables the treatment of catalogs with different sky coverage, whose angular selection function would enter the prior on the latent direction ω . Technically, the $\rho_{C_o}(\omega)$ function is not simply the angular selection function of the intersection area of the catalogs, because it also incorporates the astrometric uncertainty: there is a non-zero probability of observing a source within a given footprint even if its true direction is outside the field of view, but this effect is negligible if the field of view is large in comparison to its boundary blurred by the astrometric uncertainty, which is the case for typical observations and surveys, but would not apply, for example, if a catalog were an aggregation of disjoint sky patches comparable in size to the point-spread function.

Alternatively, one can define the marginal likelihood for a non-association hypothesis, which assumes that every source in S_o is a separate object on its own

$$\mathcal{M}_o^{\text{NA}} = \prod_{(i,c) \in S_o} \int d\omega \rho_c(\omega) \ell_{ic}(\omega), \quad (3)$$

where $\rho_c(\omega)$ is the prior probability density function of direction for sources in catalog c . This hypothesis serves as a natural comparison, with which it is useful to introduce the Bayes factor as the ratio of the marginal likelihoods of these two cases,

$$B_o = \frac{\mathcal{M}_o}{\mathcal{M}_o^{\text{NA}}}. \quad (4)$$

The B_o takes values larger than 1 when the association of the sources in S_o is more likely than the alternative, $B_o < 1$ favors separate objects. We note that $\prod \mathcal{M}_o^{\text{NA}}$ is simply a product over all sources in all catalogs independent of partition P and, hence, is constant. Consequently, the maximization of the likelihood $\mathcal{L}(P)$ is equivalent to optimizing $\prod B_o$.

As customary, we work with a summary of the raw imaging data D_{ic} for each detection (i, c) , the measured direction x_{ic} , which is essentially the intensity-weighted pixel direction. In order to calculate the Bayes factors B_o with these measurements, we specify a distribution for the member likelihood function $\ell_{ic}(\omega)$, i.e., the astrometric uncertainty. To describe directional uncertainty

in the observations, the spherical analog of the Gaussian is often assumed, the Fisher (1953) distribution,

$$\ell_{ic}(\omega) := f(x_{ic}; \omega, \kappa_{ic}) = \frac{\kappa_{ic}}{4\pi \sinh \kappa_{ic}} \exp(\kappa_{ic} \omega \cdot x_{ic}), \quad (5)$$

where the x_{ic} observed direction and true ω direction are both 3D unit vectors. The latter is the mode of the distribution, and κ_{ic} is a concentration parameter. When $\kappa_{ic} \gg 1$, the Fisher distribution approximates a Gaussian distribution with standard deviation (in radians) for each coordinate σ_{ic} with $\kappa_{ic} = 1/\sigma_{ic}^2$ and the (all-sky) Bayes factor can be calculated analytically as shown in Budavári & Szalay (2008) as follows,

$$B_o = 2^{|S_o|-1} \frac{\prod_{ic} \kappa_{ic}}{\sum_{ic} \kappa_{ic}} \exp\left(-\frac{\sum_{ic} \sum_{i'c'} \kappa_{ic} \kappa_{i'c'} \psi_{ic,i'c'}^2}{4 \sum_{ic} \kappa_{ic}}\right), \quad (6)$$

where (i, c) and (i', c') are all sources in subset S_o and $\psi_{ic,i'c'}$ is the (small) angle between the directions for sources (i, c) and (i', c') .

The next section discusses how to find the globally optimal associations, i.e., the partition P that maximizes the $\mathcal{L}(P)$ likelihood function by optimizing $\prod B_o$ using integer linear programming.

2.1. CanILP: Optimal Selection of Candidates

First we summarize the previous approach introduced by Shi et al. (2019) and highlight some of the outstanding challenges. Maximizing $\prod B_o$ is equivalent to minimizing

$$-\sum_o \ln B_o. \quad (7)$$

Given a data set D of all (i, c) pairs for all catalog c and source i in catalog c , we introduce a binary variable x_T taking values in $\{0, 1\}$ for each nonempty subset $T \subseteq D$, with the interpretation that $x_T = 1$ indicates that the subset T is included in the partition. To ensure the validity of the partition, we require

$$\sum_{T \ni (i,c)} x_T = 1 \quad (8)$$

for every element $(i, c) \in D$. This forces every source (i, c) to be included in exactly one subset of the partition. However, note that for an orphan o , $B_o = 1$. Hence, these coefficients do not contribute to the objective function and we could simply remove those subsets T that have $|T| = 1$. From this, we can modify the above constraint to

$$\sum_{T \ni (i,c)} x_T \leq 1 \quad (9)$$

for every element $(i, c) \in D$. In the final solution, if a source (i, c) does not appear in any subset T , we treat it

as an orphan. For example, in Figure 1, Source (2,1) is not included in any subset T , so, in the solution, we will include it as an orphan.

By defining $w_T = -\ln B_T$ for every subset T , the final integer linear programming function can be written as follows,

$$\begin{aligned} & \min \sum_T w_T x_T \\ & \text{subject to } x_T \in \mathbb{Z} \text{ and } 0 \leq x_T \leq 1 \text{ for all } T, \\ & \text{and } \sum_{T \ni (i,c)} x_T \leq 1 \text{ for all } (i, c) \in D. \end{aligned} \quad (10)$$

Note that the formulation above can be used to solve the matching problem given any number of catalogs C but requires an enumeration of all possible candidate associations, which can be numerous. This requires calculations of combinatorially many w_T and the introduction of the corresponding x_T binary variables, which quickly becomes prohibitively expensive for many catalogs.

Shi et al. (2019) demonstrated their approach on three catalogs with identical astrometric uncertainty. Here we extend their work and describe a novel approach, which can be efficiently used with many catalogs.

2.2. DirILP: Optimal Direct Associations

The key idea is to introduce variables that directly assign the detections to hypothesized objects instead of simply switching on and off some previously enumerated candidate associations in the final matched catalog. The objective $\prod B_o$ is the same but expressing it with the new variables is significantly more complicated than before. In the process one needs to introduce several additional (sets of) auxiliary variables to linearize the problem. In case of homoscedasticity when all astrometric uncertainty are the same for all detections, the linearization is relatively straightforward, but further modeling tricks are required in the general setting. In the following sections these two cases are introduced along with the variables needed to model and solve the global association problem. Further details are provided in the appendix about the derivation of the general heteroscedastic formalism.

2.2.1. Homoscedasticity

For simplicity, we first discuss the special case where the astrometric uncertainty of each detection is the same, i.e., $\sigma_{ic} = \sigma$ for each source (i, c) . Given a data set D , let N be the total number of detections in all catalogs considered. The number of astronomical objects these represent will be at most N , corresponding to the hypothesis that every detection comes from a different object. Our goal is to find a mapping that matches each

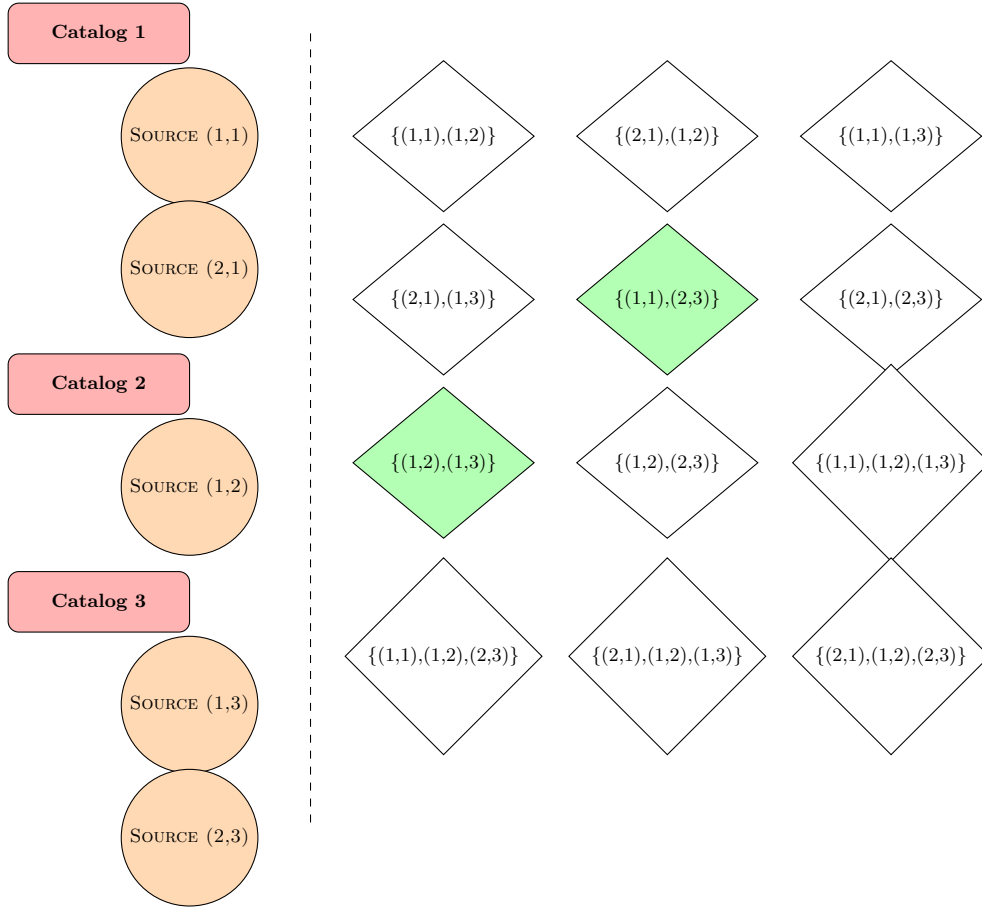


Figure 1. An illustration of CanILP. As can be seen on the left side, we assume there are 2 detections from Catalog 1 (Sources (1, 1) and (2, 1)), 1 detection from Catalog 2 (Source (1, 2)) and 2 detections from Catalog 3 (Sources (1, 3) and (2, 3)). In CanILP, we list all candidates for possible associations across independent detections, which are shown on the right side. These are the x_T in the formulation. We then find the combinations of subsets that maximize the overall likelihood. Here, the solution given by CanILP indicates that the subsets $\{(1, 1), (2, 3)\}$ and $\{(1, 2), (1, 3)\}$ are included in the partition. These subsets, which are represented by a green color, correspond to the variables $x_{\{(1,1),(2,3)\}} = x_{\{(1,2),(1,3)\}} = 1$ in the model. On the other hand, all other variables $x_T = 0$. Notice that because Source (2, 1) does not appear in any of these subsets, so we treat it as an orphan. As a result, the association outputted by CanILP is $\{(1, 1), (2, 3)\}, \{(1, 2), (1, 3)\}, \{(2, 1)\}$.

source to one (and only one) object. This association between a source and an object means that the source is an observation of that object in the sky. Naturally, multiple sources are expected to be assigned to the same object, which represents the hypothesis that all of these sources are observations of that same object. To capture the matching between a source (i, c) and an object o , we introduce binary variables $\{x_{ic}^o\}$, where a given $x_{ic}^o = 1$ if the (i, c) detection is associated with object o , and 0 otherwise.

Figure 2 illustrates how this approach works. For example, the arrow from Source (2, 1) to Object 1 representing an association means that $x_{21}^1 = 1$. Similarly, $x_{11}^3 = 0$ means no association, hence there is no arrow between the corresponding entries. A partition P can now be represented as a set $\{S_o : o = 1, \dots, N\}$, where

S_o is the subset of sources assigned to o , i.e.,

$$S_o := \{(i, c) : x_{ic}^o = 1\}. \quad (11)$$

If, for a given index o , $x_{ic}^o = 0$ for all (i, c) sources, then $S_o = \emptyset$ is empty, which means object o is not needed for that particular partition.

Recall that the goal is to maximize the product of Bayes factors $\prod B_o$ (or to minimize $-\sum \ln B_o$) corresponding to these associations. Given an association S_o , assuming $\kappa_{ic} = \kappa$ for all source (i, c) , eq. (6) gives us

$$B_o = 2^{|S_o|-1} \frac{\prod_{ic} \kappa}{\sum_{ic} \kappa} \exp\left(-\frac{\sum_{ic} \sum_{i'c'} \kappa^2 \psi_{ic,i'c'}^2}{4 \sum_{ic} \kappa}\right) \quad (12)$$

$$= 2^{|S_o|-1} \frac{\kappa^{|S_o|}}{|S_o| \kappa} \exp\left(-\frac{\kappa \sum_{ic} \sum_{i'c'} \psi_{ic,i'c'}^2}{4 |S_o|}\right) \quad (13)$$

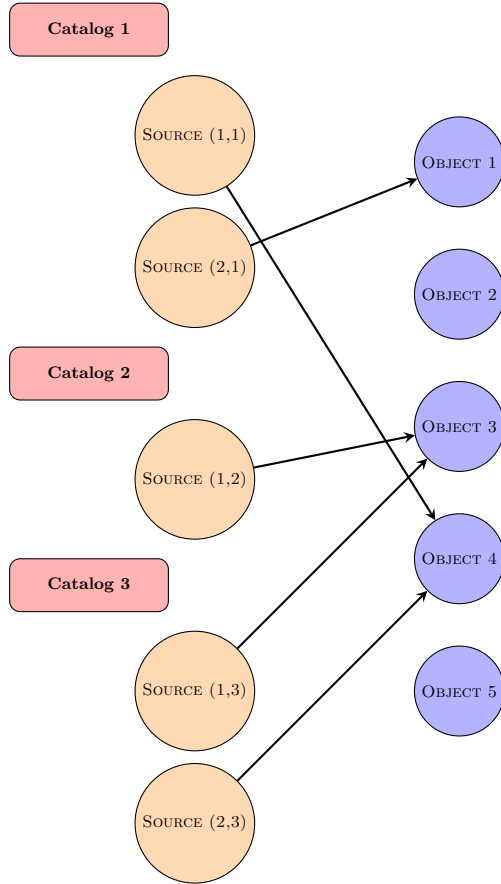


Figure 2. An illustration of DirILP. As in Figure 1, assume there are 2 detections from Catalog 1 (Sources (1,1) and (2,1)), 1 detection from Catalog 2 (Source (1,2)) and 2 detections from Catalog 3 (Sources (1,3) and (2,3)). In this case, the output of DirILP indicates that Sources (1,1) and (2,3) belong to the same object, that Sources (1,2) and (1,3) belong to the same object, and that Source (2,1) is an orphan. Notice that it is okay for an object to not have any source associated with it. The solution given by DirILP is $\{\{(1,1), (2,3)\}, \{(1,2), (1,3)\}, \{(2,1)\}\}$, which is the same as the one given by CanILP in Figure 1.

Hence,

$$-\ln B_o = \ln(2\kappa)(1 - |S_o|) + \ln |S_o| + \frac{\kappa \sum \psi_{ic,i'c'}^2}{4|S_o|} \quad (14)$$

We want to find the partition P that minimizes $-\sum \ln B_o$. Notice that as of now, there are still several non-linear terms in $-\ln B_o$ so it is not yet a linear objective. To make use of ILP method, we will first need to rewrite this as a linear function. To do that, we introduce the following variables, defined for each index $k \in \{0, \dots, C\}$, with C representing the total number of catalogs:

$$z_k^o = \begin{cases} 1 & \text{if } \sum_{ic} x_{ic}^o = k \\ 0 & \text{otherwise} \end{cases} \quad (15)$$

This variable captures the number of sources getting matched to object o , or the cardinality of the subset S_o . When $z_k^o = 1$, there are k hypothesized observations of object o in the data. In addition, notice that at most 1 of the z_k^o , $k = 0, \dots, C$ can be 1. We also introduce

$$y_{ic,i'c'}^o = \begin{cases} 1 & \text{if } x_{ic}^o = x_{i'c'}^o = 1 \\ 0 & \text{otherwise} \end{cases} \quad (16)$$

This is an indicator variable that checks whether the sources (i, c) and (i', c') belong to the same object o . In particular, $y_{ic,i'c'}^o = 1$ indicates the hypothesis that sources (i, c) and (i', c') are observations of object o . We also have

$$t^o = \begin{cases} \frac{\sum \kappa \psi_{ic,i'c'}^2 y_{ic,i'c'}^o}{4k} & \text{if } z_k^o = 1 \text{ for some } k \in [C] \\ 0 & \text{if } z_0^o = 1 \end{cases} \quad (17)$$

where $[C]$ represents the set of numbers $\{1, 2, \dots, C\}$. This variable captures the last term in $-\ln B_o$ for a subset S_o . In particular, when $z_k^o = 1$ for some $k \in [C]$, i.e. $|S_o| = k$ by definition of z_k^o , we have

$$t^o = \frac{\sum \kappa \psi_{ic,i'c'}^2}{4|S_o|} \quad (18)$$

as desired, where the summation goes over all (i, c) and (i', c') in S_o . On the other hand, when $z_0^o = 1$, no detection is assigned to object o , so this term should contribute nothing to the objective function. Next, we introduce

$$p^o = \begin{cases} (1 - k) \ln(2\kappa) & \text{if } z_k^o = 1 \text{ for some } k \in [C] \\ 0 & \text{if } z_0^o = 1 \end{cases} \quad (19)$$

This variable captures the first term in $-\ln B_o$ for a subset S_o . It plays a similar role as t^o , i.e., when $z_0^o = 1$, no detection is assigned to object o , so this term should contribute nothing to the objective function. On the other hand, if some sources are matched to object o , $p^o = \ln(2\kappa)(1 - |S_o|)$ as desired.

Finally, we will linearize the term $\ln |S_o|$ by breaking the natural log function into finitely many affine linear pieces. We first introduce constants a_1, a_2, \dots, a_C , where $a_1 = 0$ and $a_p = \ln(p) - \ln(p-1)$, for $p = 2, \dots, C$. Then for each object o , we define binary variables $w_1^o \geq w_2^o \geq \dots \geq w_C^o$ and impose the constraint that

$$\sum_{p=1}^C w_p^o = \sum_{ic} x_{ic}^o = |S_o|. \quad (20)$$

Using the new notation, we can now express $\ln |S_o|$ as a linear function of w_p^o : $\ln |S_o| = \sum_{p=1}^C a_p w_p^o$. To explain

why this is the case, it is best to work with an example. Suppose 3 sources are matched with object o , so $|S_o| = 3$ and $\ln|S_o| = \ln 3$. Because $\sum_{p=1}^C w_p^o = |S_o| = 3$ and w_p^o are 0/1 variables with $w_1^o \geq w_2^o \geq \dots \geq w_C^o$, we have $w_1^o = w_2^o = w_3^o = 1$ and $w_4^o = w_5^o = \dots = w_C^o = 0$. Then, $\sum_{p=1}^C a_p w_p^o = a_1 + a_2 + a_3 = (0) + (\ln 2 - \ln 1) + (\ln 3 - \ln 2) = \ln 3$, which is exactly $\ln|S_o|$. Our objective function now becomes

$$\min \sum_o \left(p^o + \sum_p a_p w_p^o + t^o \right), \quad (21)$$

which is linear in the variables p^o, w_p^o and t^o .

As can be seen in the definitions of these variables, there are certain relationships that still need to be modeled using linear constraints. The full ILP formulation is given in Appendix A with detailed explanations for how the constraints model the relationships between the variables $x_{ic}^o, y_{ic,i'c'}, z_k^o, p^o, w_p^o$, and t^o .

2.2.2. Heteroscedasticity

We can also remove the assumption that every source has the same measure of uncertainty κ_{ic} . From eq. (6), we have,

$$-\ln B_o = (1 - |S_o|) \ln 2 - \sum_{ic} \ln \kappa_{ic} + \ln \sum_{ic} \kappa_{ic} + \frac{\sum_{ic} \sum_{i'c'} \kappa_{ic} \kappa_{i'c'} \psi_{ic,i'c'}^2}{4 \sum_{ic} \kappa_{ic}}, \quad (22)$$

where all the summations run over all (i, c) and (i', c') in S_o . We use x_{ic}^o, z_k^o , and $y_{ic,i'c'}$ as defined in the special case of Section 2.2. We also introduce new variables to convert eq. (22) into a linear function.

We first linearize the term $\ln \sum_{ic} \kappa_{ic}$ using the same trick as when we linearized $\ln|S_o|$ in Section 2.2. We introduce constants $b_{\min} \equiv b_1, b_2, b_3, \dots$, where

$$b_{\min} = \ln \left(\min_{ic \in D} \kappa_{ic} \right) \quad (23)$$

and

$$b_{\max} = \ln \left(C \max_{ic \in D} \kappa_{ic} \right). \quad (24)$$

Now, if we set an error threshold ϵ , then the

$$R \equiv \left\lceil \frac{b_{\max} - b_{\min}}{\epsilon} \right\rceil \quad (25)$$

constants b_p are defined as

$$b_p = b_{\min} + (p - 1) \times \epsilon \quad \text{for } p = 1, \dots, R. \quad (26)$$

Then for each object o , we define binary variables $\chi_1^o \geq \chi_2^o \geq \dots \geq \chi_R^o$ and impose the constraint

$$\chi_1^o \exp(b_1) + \sum_{p=2}^R \chi_p^o [\exp(b_p) - \exp(b_{p-1})] \geq \sum_{ic} \kappa_{ic} x_{ic}^o. \quad (27)$$

Using the new variables, we have

$$\ln \sum_{ic} \kappa_{ic} \approx \chi_1^o b_1 + \sum_{p=2}^R \chi_p^o (b_p - b_{p-1}) = \chi_1^o b_{\min} + \epsilon \sum_{p=2}^R \chi_p^o \quad (28)$$

since $b_p - b_{p-1} = \epsilon$ for all $p \geq 2$.

To illustrate how the χ_p^o variables work, let us assume that after looking at the data, we determine that $b_{\min} = 29$ and $b_{\max} = 33$. If we let $\epsilon = 0.5$, then the value of constants b_p are $\{29, 29.5, \dots, 32.5, 33\}$. Now suppose there are 3 sources that are matched to an object o with associated κ_{ic} values of $5 \times 10^{12}, 8 \times 10^{12}$, and 10^{13} . Then the true value of $\ln \sum_{ic \in S_o} \kappa_{ic}$ is $\ln(2.3 \times 10^{13})$, which evaluates to 30.77. With the defined variables, the solution given by ILP is $\chi_1^o = \chi_2^o = \dots = \chi_5^o = 1$ and $\chi_6^o = \dots = \chi_9^o = 0$ because $\chi_1^o \exp(b_1) + \sum_{p=2}^P \chi_p^o (\exp(b_p) - \exp(b_{p-1})) = \exp(29) + \exp(29.5) - \exp(29) + \exp(30) - \exp(29.5) + \dots + \exp(31) - \exp(30.5) = \exp(31) > 2.3 \times 10^{13}$, which satisfies the constraint $\chi_1^o \exp(b_1) + \sum_{p=2}^P \chi_p^o (\exp(b_p) - \exp(b_{p-1})) \geq \sum_{ic} \kappa_{ic} x_{ic}^o$. Notice that setting the variables $\chi_6^o, \dots, \chi_9^o = 1$ will also satisfy the constraint. However, since we will model our problem with a minimization objective, the optimal solution will force $\chi_1^o b_1 + \sum_{p=2}^R \chi_p^o (b_p - b_{p-1})$ to be as small as possible. Finally, notice that in this case the value of $\chi_1^o b_1 + \sum_{p=2}^R \chi_p^o (b_p - b_{p-1})$, which is used to approximate $\ln \sum_{ic \in S_o} \kappa_{ic}$, is 31, which is close to the true value of 30.77.

Next, we will linearize the last term in eq. (22) by first introducing the constant

$$c_{\min} = \min_{ic \in D} \kappa_{ic} \quad (29)$$

and

$$c_{\max} = C \max_{ic \in D} \kappa_{ic}. \quad (30)$$

Then by rounding these two values to the nearest 100, we can introduce grid points $0 \equiv c_0, c_1, c_2, \dots, c_Q$, where c_1 is the nearest 100 of c_{\min} , c_Q is the nearest 100 of c_{\max} , and for all $i > 2$, $c_i = c_1 + 100(i - 1)$. We then introduce

$$u_k^o = \begin{cases} 1 & \text{if } \sum_{ic} [\kappa_{ic}]_{100} x_{ic}^o = c_k \\ 0 & \text{otherwise} \end{cases} \quad (31)$$

where k ranges in $\{0, 1, \dots, Q\}$ and the operator $[\cdot]_{100}$ is defined as rounding to the nearest 100. This variable attempts to approximate $\sum_{ic \in S_o} \kappa_{ic}$, which appears in the denominator of the last term of eq. (22).

The variables p^o and t^o are also very similar to the definitions in Section 2.2; however, we need to slightly modify them as follows:

$$t^o = \frac{\sum_{ic} \sum_{i'c'} \kappa_{ic} \kappa_{i'c'} \psi_{ic,i'c'}^2 y_{ic,i'c'}^o}{4c_k}, \quad (32)$$

if $u_k^o = 1$ for some $k \in \{1, 2, \dots, Q\}$, and $t^o = 0$ otherwise.

The reasoning for defining t^o this way is that if $u_0^o = 1$,

$$\sum_{(i,c)} [\kappa_{ic}]_{100} x_{ic}^o = c_0 = 0. \quad (33)$$

This happens only when $x_{ic}^o = 0$ for all (i, c) , i.e. no sources are matched to object o . Hence, t^o should not contribute to the objective function, hence the value of 0. On the other hand, if $u_k^o = 1$ for some $k > 0$, by definition of u_k, c_k is the best approximation to $\sum_{ic \in S_o} \kappa_{ic}$. Thus, (32) holds.

In addition, we modify p^o defined in eq. (19) as follows,

$$p^o = \begin{cases} (1-k) \ln(2) & \text{if } z_k^o = 1 \text{ for some } k \in [C] \\ 0 & \text{if } z_0^o = 1 \end{cases}. \quad (34)$$

This variable serves a similar function as in the homoscedastic case, which is to capture the first term in eq. (22).

The objective function can now be written as

$$\sum_o \left(p^o - \sum_{ic} x_{ic}^o \ln \kappa_{ic} + \chi_1 b_{\min} + \epsilon \sum_{p=2}^P \chi_p^o + t^o \right), \quad (35)$$

which is linear in all the variables involved.

There are certain relationships that still need to be modeled using linear constraints because ILP formulations only take linear constraints. The full ILP formulation is given in Appendix B, with detailed explanations for how the constraints model the relationships between the variables $x_{ic}^o, y_{ic,i'c'}, z_k^o, \chi_p^o, u_k^o, p^o$, and t^o .

3. MOCK OBJECTS AND SIMULATIONS

We consider the idealized case where all the catalogs capture the same astronomical properties of objects in the sky, i.e., they detect the same set of objects. As we generate 100 objects and assume there are C distinct catalogs, we expect to see $100 \times C$ sources and 100 C -way association sets. We will now show the catalog matching results using both of our approaches. The ILP programs in both approaches are solved using Gurobi, an optimization solver [Gurobi Optimization \(2020\)](#).

3.1. Homoscedasticity: identical $\kappa_{ic} = 1/\sigma^2$

Observe that for the CanILP formulation in Section 2.1, we need to list all the possible valid subsets $T \subseteq D$. We could do this by sequentially adding catalogs one by one and considering sources from the new catalog. However, this evaluates to $101^C - 1$ subsets, which is exponential in terms of the number of catalogs C . Hence,

we first try to reduce the number of possible subsets by observing that sources that are far away cannot be from the same object. So we can impose some distance constraints on the sources that are put into the same candidate association set. In doing so, we should be careful not to discard potentially true associations later on because say two sources from the first 2 catalogs that are far away might not be a 2-way matching; but, if on the third catalog, there is a source lying in the middle of the path between these 2 sources, the 3 sources together might be a 3-way matching.

That being said, this suggests an idea for dividing the whole region of the sky that is of interest into different islands where the sources are clustered together so that instead of solving one big problem, we could break it into smaller problems and make use of parallel computing. Essentially, we first apply a single-linkage clustering algorithm to our dataset, which is done using the DBSCAN algorithm with parameters “min_samples” = 2 and “eps” = $5 \times \max_{ic \in D} \sigma_{ic}$. With the chosen parameters, we are essentially performing the friends-of-friends algorithm. It turns out that for our simulation, most of these islands consist of only 1 source from each catalog. Hence, from now on, we will show the result for this scenario of having 1 source from each catalog. This situation is not peculiar to our simulation but is, in fact, observed in real data sets from multiple visits of the same part of the sky by the same telescope. Analysis for the multiple sources per catalog will be discussed later. As can be seen in Figure 3, even though we are able to handle more than 3 catalogs, the maximum number of catalogs we could analyze in a day is 20. Though, similar to [Shi et al. \(2019\)](#), we do not include any pruning procedures, such as those described in [Kunszt et al. \(2001\)](#); [Gorski et al. \(2005\)](#); [Gray et al. \(2007\)](#); [Lee & Budavári \(2017\)](#). These pruning procedures might speed up the matching, but from our experience the complexity of the problem is still exponential in terms of the number of catalogs. The next paragraph discusses how far we could get using DirILP formulation.

DirILP formulation analysis.—The main drawback from the previous approach is that the process of creating potential subsets T is exponential in terms of the number of catalogs. Even if we consider the island, the number of nonempty subsets in such an island will still be $2^C - 1$, so creating the variables for the ILP takes a tremendous amount of time.

DirILP formulation attempts to fix that problem by reducing the time complexity to create the variables for the ILP to something that is polynomial in the total number of sources. However, since this catalog matching problem is intrinsically difficult, we still have to tackle

the exponential complexity of the problem somewhere else: this appears in the time needed to solve the ILP. We believe that with advances in the field of integer linear programming, the Gurobi solver will be able to solve this problem more efficiently. It turns out using DirILP, we are able to tackle up to 60 catalogs. The comparison for the total running time between CanILP and DirILP is shown in Figure 3. In addition, we also include the set up time and optimization time for each formulation in Figures 4 and 5.




Figure 3. Total running time comparison between the two formulations for the special case (Log Scale). Notice that CanILP chokes when there are 20 catalogs.

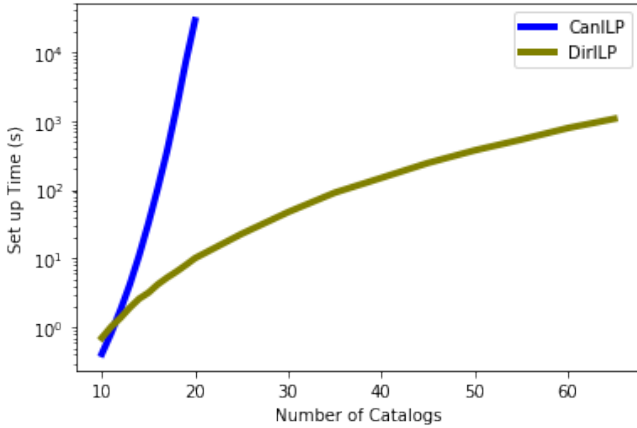


Figure 4. Set up time comparison between the two formulations for the special case (Log Scale)

Moreover, by including some heuristic constraints, such as imposing a time limit between incumbent solutions, on the Gurobi solver, we are able to push the DirILP further to handle 160 catalogs.

Finally, it is important to note that the associations given by CanILP and DirILP are the same and they

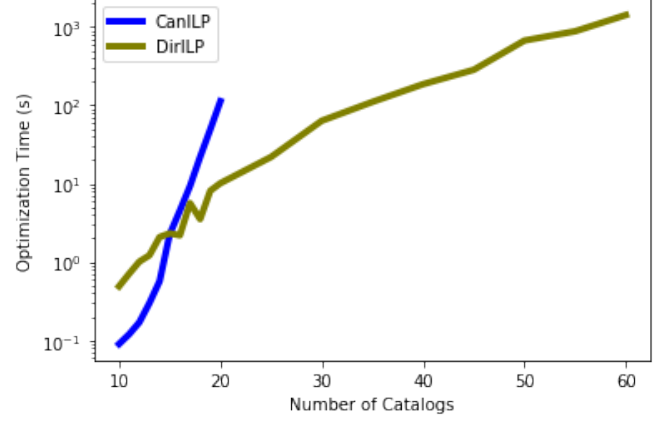


Figure 5. Optimization time comparison between the two formulations for the special case (Log Scale)

match the ground truth perfectly. Hence, there is no difference in the accuracy of the matching between the two approaches. They only differ in their running time.

3.2. General case: different κ_{ic} for every detection

For the general case, both approaches still give all correct associations that match the ground truth. However, as in the special case, DirILP is still more efficient at solving the matching problem than CanILP, as shown in Figure 6. We should point out that even though in this general setting, the optimal value found in DirILP is just an approximation of the Bayes factor associated with the ground-truth matching, the values are still quite close to each other. More importantly, the associations obtained from DirILP still match the ground-truth associations. Figures 6 - 8 show the total running time, time to set up the ILP, and time for Gurobi to solve the ILP, for both CanILP and DirILP in this general case.




Figure 6. Total running time comparison between the two formulations for the general case (Log Scale). Notice that CanILP chokes when there are 18 catalogs.




Figure 7. Set up time comparison between the two formulations for the general case (Log Scale)

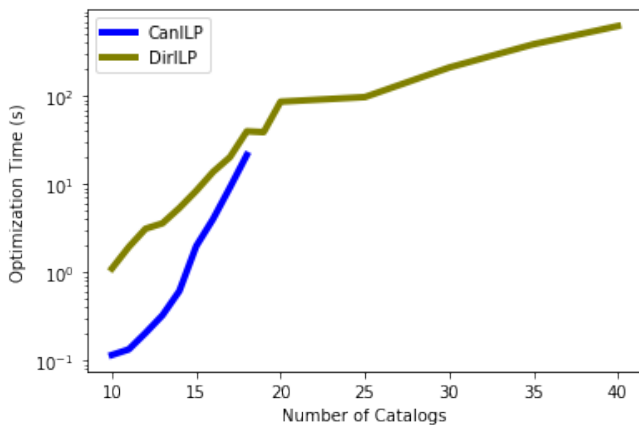


Figure 8. Optimization time comparison between the two formulations for the general case (Log Scale)

3.3. Multiple sources per catalog in each island

Recall that in the previous sections, we assume that in each island there is only one detection from each catalog, which is a reasonable assumption in many real-life situations. In this section, we would like to discuss scenarios when the uncertainty σ_{ic} is large or the source density is very high. These scenarios will result in islands where there might be multiple detections from each catalog in an island. It turns out that in our simulation, CanILP and DirILP still give the correct association under this scenario. However, both methods run much slower than in the previous scenario and are able to handle only about half as many catalogs with the same settings on the algorithms. We give an example of the running time for the 2 methods when there are 2 detections from each catalog. One can see how much worse it can get when the number of detections from

each catalog becomes larger. Figure 9 shows the total running time for both CanILP and DirILP when there are 2 detections from each catalog in an island.




Figure 9. Total Running time comparison between the two formulations when there are 2 detections from each catalog in an island (Log Scale)

3.4. Discussion of running time complexity

We now give a brief explanation for the shape of the curves in Figures 3–8. For CanILP, since the number of variables is exponential in terms of the number of catalogs, under the log scale as in Figures 4 and 7, the time to create these variables and set up the ILP as a function of the number of catalogs is represented by a straight line. On the other hand, for DirILP, we have a curve with decreasing gradient instead of a straight line because the number of variables and constraints in this method is polynomial in the number of catalogs. The explanation for the curves in Figures 5 and 8 are similar because the amount of time to solve an ILP generally depends on the number of variables and constraints in the problem. That being said, the curves in these two figures look more jagged because of other complexities involved in the optimization procedure. Finally, as most of the time to solve the catalog matching problem is spent on setting up the ILP, the curves in Figures 3 and 6 are very much similar to their counterparts in Figures 4 and 7, respectively.

4. IMPLEMENTATION AND SOFTWARE

CanILP and DirILP algorithms are implemented in several Jupyter notebooks. They share a common structure: In the first part, we create a simulation with different catalogs and a number of mock objects on each of the catalogs. Next, we perform the DBSCAN algorithm to output different islands, or clusters of detections. Again, as mentioned in 3.1, with our chosen parameters, this is

similar to executing a friends-of-friends algorithm. The reason we pick DBSCAN is because of its well-developed library in Python.

After running the clustering algorithm, we implement CanILP and DirILP to solve the catalog matching problem in each island. The optimization problems in these modules were solved using Gurobi software [Gurobi Optimization \(2020\)](#).

In addition, we employ several (optional) heuristics in the DirILP algorithm to speed up the catalog matching procedure. First, any 2 sources that are more than 8σ away from each other, we force them to belong to separate objects. Second, for sources that are 0.1σ away from each other, we restrict them to belong to the same object. Finally, we set an MIP Gap (optimality gap) of 0.5% to prevent Gurobi from taking too long to verify optimality. Through our experiments, we have found that running the algorithm with these heuristics give the same results but it was 10 times faster. The notebooks can be found on Github at <https://github.com/tunguyen52/Nway-matching>.

5. SUMMARY AND FUTURE WORK

We have shown how the CanILP approach of [Shi et al. \(2019\)](#) and the new DirILP solve for a globally optimal matched catalog in crowded fields where a greedy approach is not sufficient. The former enumerates the candidate associations and picks the optimal combination of those; the latter introduces variables to directly assign sources to objects. The new DirILP formulation is superior in the sense that it scales to large number of catalogs better, i.e., produces results in less time. The method comes at a price, which is in complexity of the algorithm, especially in case of heteroscedasticity when the catalogs have different astrometric uncertainty.

In fact DirILP only out-performs the previous method when many catalogs are to be crossmatched. We recommend the simpler CanILP approach for small number of catalogs, where the combinatorial explosion is not as severe. In our experiments, this crossover threshold appears to be at around 12 visits or catalogs, beyond which the DirILP method gets faster.

Both of these methods optimize the same objective and yield the best possible catalog matching result in a likelihood sense. No prior on the partition is imposed currently in our study and the accompanying software. While placing priors on the partition might seem complicated, certain simple priors can be easily expressed, such as those that depend only the number of objects in the matched catalog. This is a possible direction to explore in the future.

Additional future work includes testing the software and its performance on imaging surveys, such as multiple visits of the Hyper Suprime-Cam Subaru Strategic Program, whose data collection resembles future observations of the LSST.

- 1 This material is based upon work supported by the Na-
- 2 tional Science Foundation under Grant No. 1909709,
- 3 1814778, 1452820 and 1934979. T.B. gratefully acknowl-
- 4 edges the aforementioned AST/AAG grants from NSF
- 5 and funding from NASA via awards STScI-49721 and
- 6 STScI-52333 under NAS5-26555. T.N. and A.B. grate-
- 7 fully acknowledge support from the aforementioned NSF
- 8 grant and ONR grant N000141812096. A.B. is also
- 9 grateful for support from the aforementioned NSF grant
- 10 and AFOSR grant FA95502010341. The authors thank
- 11 the anonymous referee for the careful review and the
- 12 thoughtful comments.

APPENDIX

A. DIRILP FORMULATION - SPECIAL CASE

The ILP Formulation for DirILP when $\kappa_{ic} = \frac{1}{\sigma^2}$ for every source (i, c) is given below. The objective function we want to minimize is given by

$$\min \sum_o \left(p^o + \sum_p a_p w_p^o + t^o \right) \quad (\text{A1})$$

The following constraints restrict $x_{ic}^o, y_{ic,i'c'}^o, z_k^o, w_p^o$ to binary variables and t^o to have non-negative values.

$$x_{ic}^o, y_{ic,i'c'}^o, z_k^o, w_p^o \in \mathbb{Z} \text{ and } 0 \leq x_{ic}^o, y_{ic,i'c'}^o, z_k^o, w_p^o \leq 1, 0 \leq t^o, \quad \forall(i, c), k, p, o. \quad (\text{A2})$$

The next equation ensures that all sources (i, c) need to belong to exactly one subset:

$$\sum_o x_{ic}^o = 1, \quad \forall(i, c) \quad (\text{A3})$$

The following equation imposes that every subset takes no more than 1 source from each catalog.

$$\sum_i x_{ic}^o \leq 1, \quad \forall o \in \{1, 2, \dots, N\}, \quad \forall c \in \{1, \dots, C\} \quad (\text{A4})$$

The following set of constraints on $y_{ic,i'c'}^o$ is an implementation of the definition of $y_{ic,i'c'}^o$ in Section 2.2, which requires $y_{ic,i'c'}^o = 1$ only if $x_{ic}^o = x_{i'c'}^o = 1$:

$$y_{ic,i'c'}^o \geq x_{ic}^o + x_{i'c'}^o - 1, \quad (\text{A5})$$

$$y_{ic,i'c'}^o \leq x_{ic}^o, \quad (\text{A6})$$

$$y_{ic,i'c'}^o \leq x_{i'c'}^o, \quad (\text{A7})$$

for all $(i, c) \neq (i', c')$ and $\forall o$.

Since the cardinality of any subset from a partition P is between 0 and C , the following equation states that only 1 of z_k^o can take a value of 1.

$$\sum_{k=0}^C z_k^o = 1, \quad \forall o, \quad (\text{A8})$$

The next constraint is the definition of w_p^o as described in Section 2.2.

$$w_1^o \geq w_2^o \geq \dots \geq w_C^o \quad \text{and} \quad \sum_{p=1}^C w_p^o = \sum_{ic} x_{ic}^o, \quad \forall o, \quad (\text{A9})$$

With the specific choice of the constant M as defined below, the equation that follows becomes redundant when $z_k^o = 0$ since the RHS will be negative and so $t_o \geq 0$ becomes the enforcing constraint, and when $z_k^o = 1$, the minimization forces t^o to be equal to the first term of the RHS.

$$t^o \geq \frac{\sum \kappa \psi_{ic,i'c'}^2 y_{ic,i'c'}^o}{4k} - (1 - z_k^o)M, \quad \forall o \text{ and } k \in \{1, 2, \dots, C\}, \quad (\text{A10})$$

where $M = \left\lceil \sum_{ic,i'c' \in D} \frac{\kappa \psi_{ic,i'c'}^2}{4} \right\rceil$.

The following set of equations constitutes the definition of z_k^o .

$$\sum_{ic} x_{ic}^o \leq k z_k^o + C(1 - z_k^o) \quad (\text{A11})$$

$$\sum_{ic} x_{ic}^o \geq k z_k^o, \quad (\text{A12})$$

for all $k \in \{0, 1, 2, \dots, C\}$ and for all o .

Finally, the last equation

$$p^o \geq \ln(2\kappa)(1 - \sum_p w_p^o) - \ln(2\kappa)z_0^o, \quad \forall o, \quad (\text{A13})$$

ensures that for an empty subset S_o , $p^o = 0$, hence contributing nothing to the objective. This is because when $z_0^o = 1$ (nothing is assigned to subset S_o), $w_p^o = 0, \forall p$. As we are minimizing the objective function with respect to p^o , p^o will be set to 0. On the other hand, when $z_0^o = 0$, the constraint becomes $p^o \geq \ln(2\kappa)(1 - \sum_p w_p^o)$ and again, since we are minimizing, p^o will equal this value.

B. DIRILP FORMULATION - GENERAL CASE

Below, we give the ILP Formulation for DirILP when κ_{ic} is different for distinct sources (i, c) . Some of these constraints are similar to the special case so we will only give explanations for the new constraints, which are shown after the ILP formulation. We follow the notation introduced in Section 2.2.2. In particular, recall the constants c_0, c_1, \dots, c_Q designed to model, up to the nearest 100, the sum of subsets of uncertainties κ_{ic} , the associated decision variables u_k^o and the decision variables χ_P^o to model the logarithms of such sums.

The objective in the general case is

$$\min \sum_o \left(p^o - \sum_{ic} x_{ic}^o \ln \kappa_{ic} + \chi_1^o b_{\min} + \epsilon \sum_{p=2}^R \chi_p^o + t^o \right) \quad (\text{B14})$$

As in the special case, the following constraints on the variables restrict $x_{ic}^o, y_{ic,i'c'}^o, z_k^o$ to binary variables and t^o to have non-negative values. Additionally, the new variables χ_p^o, u_k^o are also restricted to binary values.

$$x_{ic}^o, y_{ic,i'c'}^o, z_k^o, \chi_p^o, u_k^o \in \mathbb{Z} \text{ and } 0 \leq x_{ic}^o, y_{ic,i'c'}^o, z_k^o, \chi_p^o, u_k^o \leq 1, \quad 0 \leq t^o.$$

Next, constraints (A3)–(A8) and (A11)–(A12) are included verbatim.

The following impose the conditions required on χ_p^o as described in Section 2.2.2.

$$\chi_1^o \geq \chi_2^o \geq \dots \geq \chi_R^o \text{ and } \chi_1^o \exp(b_1) + \sum_{p=2}^R \chi_p^o (\exp(b_p) - \exp(b_{p-1})) \geq \sum_{ic} \kappa_{ic} x_{ic}^o, \quad \forall o. \quad (\text{B15})$$

The next set of constraints ensure that the value of $\sum_{ic \in S_o} \kappa_{ic}$ will be approximately equal to c_k (up to the nearest 100) for some $k \in \{0, 1, 2, \dots, Q\}$.

$$\sum_{k=0}^Q u_k^o = 1, \quad \forall o, \quad (\text{B16})$$

$$\begin{aligned} \sum_{ic} [\kappa_{ic}]_{100} x_{ic}^o &\leq c_k u_k^o + M'(1 - u_k^o), \quad \forall k \in \{0, 1, 2, \dots, Q\} \text{ and } \forall o. \\ \sum_{ic} [\kappa_{ic}]_{100} x_{ic}^o &\geq c_k u_k^o \end{aligned} \quad (\text{B17})$$

where $M' = C \max_{ic \in D} \kappa_{ic}$.

Finally, the last set of constraints tie everything back into the objective function (B14).

$$t^o \geq \frac{\sum_{ic} \sum_{i'c'} \kappa_{ic} \kappa_{i'c'} \psi_{ic,i'c'}^2}{4c_k} - (1 - u_k^o)M, \quad \forall o \text{ and } k \in \{1, 2, \dots, Q\}, \quad (\text{B18})$$

$$p^o \geq \left(1 - \sum_{ic \in S_o} x_{ic}^o\right) \ln 2 - z_0^o \ln 2, \quad \forall o. \quad (\text{B19})$$

$$\text{where } M = \left\lceil \frac{\max_{ic \in D} \kappa_{ic}^2 \sum_{ic,i'c' \in D} \psi_{ic,i'c'}^2}{4 \min_{ic \in D} \kappa_{ic}} \right\rceil.$$

REFERENCES

- Boch, T., Pineau, F., & Derriere, S. 2012, in *Astronomical Society of the Pacific Conference Series*, Vol. 461, *Astronomical Data Analysis Software and Systems XXI*, ed. P. Ballester, D. Egret, & N. P. F. Lorente, 291
- Budavári, T., & Basu, A. 2016, *The Astronomical Journal*, 152, 86, doi: [10.3847/0004-6256/152/4/86](https://doi.org/10.3847/0004-6256/152/4/86)
- Budavári, T., Dobos, L., & Szalay, A. S. 2013, *Computing in Science Engineering*, 15, 12, doi: [10.1109/MCSE.2013.41](https://doi.org/10.1109/MCSE.2013.41)
- Budavári, T., & Loredó, T. J. 2015, *Annual Review of Statistics and Its Application*, 2, 113–139, doi: [10.1146/annurev-statistics-010814-020231](https://doi.org/10.1146/annurev-statistics-010814-020231)
- Budavári, T., & Szalay, A. S. 2008, *The Astrophysical Journal*, 679, 301–309, doi: [10.1086/587156](https://doi.org/10.1086/587156)
- Fisher, R. 1953, *Proceedings of the Royal Society of London Series A*, doi: [10.1098/rspa.1953.0064](https://doi.org/10.1098/rspa.1953.0064)
- Gorski, K. M., Hivon, E., Banday, A. J., et al. 2005, *The Astrophysical Journal*, 622, 759, doi: [10.1086/427976](https://doi.org/10.1086/427976)

- Gray, J., Nieto-Santisteban, M. A., & Szalay, A. S. 2007, CoRR, abs/cs/0701171.
<https://arxiv.org/abs/cs/0701171>
- Gurobi Optimization, L. 2020, Gurobi Optimizer Reference Manual. <http://www.gurobi.com>
- Kunszt, P. Z., Szalay, A. S., & Thakar, A. R. 2001, in Mining the Sky, ed. A. J. Banday, S. Zaroubi, & M. Bartelmann (Berlin, Heidelberg: Springer Berlin Heidelberg), 631–637
- Lee, M., & Budavári, T. 2017, Astronomy and Computing, 20, 155, doi: <https://doi.org/10.1016/j.ascom.2017.08.001>
- Munkres, J. 1957, Journal of the Society for Industrial and Applied Mathematics, 5
- Pineau, F.-X., Motch, C., Carrera, F., et al. 2011, Astronomy & Astrophysics, 527, A126, doi: [10.1051/0004-6361/201015141](https://doi.org/10.1051/0004-6361/201015141)
- Pineau, F.-X., Derriere, S., Motch, C., et al. 2017, Astronomy & Astrophysics, 597, A89, doi: [10.1051/0004-6361/201629219](https://doi.org/10.1051/0004-6361/201629219)
- Shi, X., Budavári, T., & Basu, A. 2019, The Astrophysical Journal, 870, 51, doi: [10.3847/1538-4357/aaf00a](https://doi.org/10.3847/1538-4357/aaf00a)
- Taghizadeh-Popp, M., Kim, J., Lemson, G., et al. 2020, Astronomy and Computing, 33, 100412, doi: <https://doi.org/10.1016/j.ascom.2020.100412>
- Taylor, M. B. 2005, in Astronomical Society of the Pacific Conference Series, Vol. 347, Astronomical Data Analysis Software and Systems XIV, ed. P. Shopbell, M. Britton, & R. Ebert, 29