

# PROCEEDINGS OF SPIE

[SPIDigitalLibrary.org/conference-proceedings-of-spie](https://SPIDigitalLibrary.org/conference-proceedings-of-spie)

## Optimisation of the WEAVE target assignment algorithm

Sarah Hughes, Gavin Dalton, Daniel Smith, Kenneth Duncan, David Terrett, et al.

Sarah Hughes, Gavin Dalton, Daniel Smith, Kenneth Duncan, David Terrett, Don Carlos Abrams, J. Alfonso Aguerri, Marc Balcells, Georgia Bishop, Piercarlo Bonifacio, Esperanza Carrasco, Shoko Jin, Ian Lewis, Scott Trager, Antonella Vallenari, "Optimisation of the WEAVE target assignment algorithm," Proc. SPIE 12184, Ground-based and Airborne Instrumentation for Astronomy IX, 121846I (29 August 2022); doi: 10.1117/12.2627114

**SPIE.**

Event: SPIE Astronomical Telescopes + Instrumentation, 2022, Montréal, Québec, Canada

# Optimisation of the WEAVE target assignment algorithm

Sarah Hughes<sup>a</sup>, Gavin Dalton<sup>a,b</sup>, Daniel Smith<sup>c</sup>, Kenneth Duncan<sup>d</sup>, David Terrett<sup>b</sup>, Don Carlos Abrams<sup>e</sup>, J. Alfonso Aguerrí<sup>f</sup>, Marc Balcells<sup>e,f,g</sup>, Georgia Bishop<sup>b</sup>, Piercarlo Bonifacio<sup>h</sup>, Esperansa Carrasco<sup>i</sup>, Shoko Jin<sup>a,j,k</sup>, Ian Lewis<sup>a</sup>, Scott Trager<sup>j</sup>, and Antonella Vallenari<sup>l</sup>

<sup>a</sup>Dept. of Physics, Keble Road, University of Oxford, OX1 3RH, UK

<sup>b</sup>RAL Space, Science and Technology Facilities Council, Rutherford Appleton Laboratory, Harwell Oxford, OX11 0QX, UK

<sup>c</sup>Centre for Astrophysics Research, University of Hertfordshire, College Ln, Hatfield, Hertfordshire, AL10 9AB

<sup>d</sup>Institute for Astronomy, Royal Observatory, University of Edinburgh, Blackford Hill, Edinburgh, EH9 3HJ, UK

<sup>e</sup>Isaac Newton Group, 38700 Santa Cruz de La Palma, Spain

<sup>f</sup>Instituto de Astrofísica de Canarias, 38200 La Laguna, Tenerife, Spain

<sup>g</sup>Dep. de Astrofísica, Universidad de La Laguna, 38200 La Laguna, Spain

<sup>h</sup>GEPI, Observatoire de Paris, Université PSL, CNRS, Place Jules Janssen, 92195 Meudon, France

<sup>i</sup>Instituto Nacional de Astrofísica, Óptica y Electrónica (INAOE), Mexico

<sup>j</sup>Kapteyn Instituut, Rijksuniversiteit Groningen, Postbus 800, 9700 AV Groningen, Netherlands

<sup>k</sup>SRON Netherlands Institute for Space Research, Niels Bohrweg 4, 2333 CA Leiden, Netherlands

<sup>l</sup>Osservatorio Astronomico di Padova, INAF, Vicolo Osservatorio 5, 35122, Padova, Italy

## ABSTRACT

WEAVE is the new wide-field spectroscopic facility for the prime focus of the William Herschel Telescope in La Palma, Spain. Its fibre positioner is essential for the accurate placement of the spectrograph's  $\sim 960$ -fibre multiplex. To maximise the assignment of its optical fibres, WEAVE uses a simulated annealing algorithm called Configure,<sup>1</sup> which allocates the fibres to targets in the field of view. We have conducted an analysis of the algorithm's behaviour using a subset of mid-tier WL<sup>2</sup> fields, and adjusted the priority assignment algorithm to optimise the total fibres assigned per field, and the assignment of fibres to the higher priority science targets. The output distributions have been examined, to investigate the implications for the WEAVE science teams.

**Keywords:** WEAVE, WL, simulated annealing, target assignment, optimisation, field configurations

## 1. INTRODUCTION

The William Herschel Telescope (WHT) Enhanced Area Velocity Explorer (WEAVE) is the new multi-object spectrograph (MOS) for the 4.2m WHT in La Palma, with a multiplex of up to 960 optical fibres in the MOS-mode that can be deployed across its 2 degree field of view. WEAVE has three observing modes; the individual multi-object spectrograph (MOS) fibres (960 and 940 fibres for plate A and plate B respectively), 20 mini IFU's (mIFU), and one large IFU (LIFU). Plate B has a reduced MOS fibre multiplex due to the presence of the mIFU's. A detailed description of the WEAVE facilities can be found in Dalton et al.<sup>3,4</sup> where the current integration progress is outlined in Hughes et al.<sup>5</sup> The design of WEAVE closely follows from its predecessor 2dF at the Anglo-Australian Telescope,<sup>6</sup> with two field plates on either side of a rotating tumbler system. The complete

---

Send correspondence to Sarah Hughes  
E-mail: sarah.hughes@physics.ox.ac.uk



Figure 1. The assembled top end of WEAVE on the William Herschel Telescope in La Palma, Spain. This photo was taken on the 25<sup>th</sup> of May 2022.

WEAVE prime-focus assembly, which is due to begin its commissioning period by mid-2022, is shown in figure 1. An overview of the WEAVE 5-year survey programs will be presented in Jin et al. (2022, *in preparation*).<sup>7</sup> Installed around the edge of the tumbler are the 168 retractor pulley units, which each contain six 1.3" optical fibres, organised into three tiers. This arrangement is shown in figure 3. The fibres are carefully placed in the focal plane on an invar field plate by either of the two positioning robots, which use a gantry system to move across the top of the positioner. This setup enables continuous observation whilst the next field is being prepared, with an approximate 1-hour time frame used for each MOS observation. This provides a good match between the configuration and expected observation time.

There are many rules in place that determine where a fibre can be placed in the field. They allow us to protect the fibres from being damaged, and allow us to achieve the configuration requirements with the full multiplex of fibres available. A program called Configure was written to achieve these goals. Configure uses simulated annealing<sup>8</sup> to assign each fibre to a target in the field of view. Its efficiency is crucial for the scientific output of WEAVE.




Figure 2. A schematic diagram of the positioner, identifying its main components.

WEAVE-LOFAR<sup>2</sup> (WL) is one of the largest surveys that will be conducted with WEAVE, and will be the primary source of spectroscopic information for the LOFAR (Low Frequency Array radio telescope)<sup>9</sup> key science project.<sup>2</sup> The WL MOS survey will be conducted over three tiers to efficiently sample the redshift and luminosity plane. These are, the wide survey, which will observe up to 10,000 deg<sup>2</sup> of the northern hemisphere, and selected targets to have a flux  $\geq 10$  mJy at 150 MHz; the mid-tier survey, which will observe up to 1,000 deg<sup>2</sup> and has a selection flux of  $\geq 1$  mJy at 150 MHz; and finally the deep-tier survey, which will observe  $\sim 60$  deg<sup>2</sup> for targets with a flux  $\geq 100$   $\mu$ Jy at 150 MHz. Further information on the WL survey can be found in Smith et al.<sup>2</sup> This analysis of Configure has been conducted with a subset of 34 WL fields, all from the mid-tier survey. For WL, it is important to maximise the number of high priority science targets observed in each field, however we must also utilise every fibre available for observations to effectively use WEAVE's capabilities. Our aim is to improve our understanding of Configure's behaviour, and find adaptations and affects which are not only applicable to WL, but to all other WEAVE science surveys.

In this paper, we investigate a range of input properties for Configure, to maximise the fibre multiplex used per observation and investigate the target priority assignment process. We are trying to optimise both the output for the higher priority target assignment and the total number of fibres used per field. This is for individual and overlapping fields of view. We wish to answer the following questions; 1) Does Configure distribute the fibres evenly across the field of view or do we see structures in the assignment? 2) How does changing the target priority impact the final assignment? 3) Is the nominal usage case the most effective method of using Configure? If not, then what is? 4) Are there any selection affects that need to be accounted for?




Figure 3. (*left*) The resting positions of several MOS fibres, arranged into three tiers, with two fibres sitting on each retractor porch. These retractors surround the edge of the field plate. (*right*) A configured WEAVE observational field. Each fibre has been individually moved by one of WEAVE's two robots and placed in position on the plate, to an accuracy of  $\pm 8\mu\text{m}$ .

## 2. CONFIGURE

The Configure program<sup>1</sup> was developed as part of the core WEAVE software system, based on the simulated annealing algorithm<sup>10</sup> used for the later developments of 2dF, and included some dynamic programming elements of the original 2dF algorithm developed in Oxford.<sup>6</sup>

### 2.1 Simulated annealing

Simulated annealing is one of several computational methods that can be effective for combinatorial optimisation problems e.g. The Travelling salesman.<sup>11</sup> It is a heuristic method that uses the connection to statistical mechanics for systems with many degrees of freedom to provide a framework for obtaining optimised results. It is very efficient at finding a minimum solution for a multi-parameter function within a modest time-period, using commercial computing power. It is worth noting that the final minimum solution produced may not be the




Figure 4. The fibre placements for a single observation field that has been allocated by Configure. This field has 1400 targets that are randomly distributed over the field plate.

global minimum of the system, but a local one instead. The definition of a combinatorial optimisation problem is to find an optimal object from a set of finite discrete objects.<sup>11</sup> They are NP-complete (nondeterministic polynomial time complete), meaning that while it is possible to use the brute-force search method, the time taken to find the ideal configuration increases rapidly with the number of variables. In the case of the WEAVE fibre positioner, a single field has  $\sim 1400$  potential targets, and a maximum of 960 fibres available for assignment. This leads to  $\sim 6.9 \times 10^{376}$  possible combinations, prior to simplification.

The basis of simulated annealing is the Metropolis algorithm.<sup>12</sup> It is analogous to the process used in condensed matter physics for growing uniform crystal structures. In its simplest form, for every step taken, an object (say an atom) is given a small random displacement. The resulting energy change  $\Delta E$  is calculated, and if  $\Delta E$  is negative then this displacement is accepted and becomes the starting point for the next step; otherwise the acceptance is determined probabilistically using

$$P(\Delta E) = \exp(-\Delta E/k_B T), \quad (1)$$

where  $P(\Delta E)$  is the probability of acceptance,  $k_B$  is the Boltzmann constant, and  $T$  is the effective temperature for optimisation. In practice, the temperature parameter is in the same units as the cost function,<sup>13</sup> which is the error estimate for the current model. This process simulates the random thermal motion of particles, and progresses by lowering the temperature of the system in stages. These cooling stages freeze certain objects of the system into place until no further changes can occur and a final minimum is given. An advantage of the annealing process is that it prevents extreme deviations from thermal equilibrium, preventing the formation of only locally optimal structures.<sup>8</sup>

For the target allocation process, a random solution is initially chosen that follows the assignment rules. Then a series of fibre swaps begins between pairs of targets and fibres, reducing the energy of the system which gradually “freezes” the fibres in place until a final cooling temperature is reached.

## 2.2 Phases of assignment

To save computational time, Configure begins by calculating a collision matrix, which records whether there is a conflict for every possible pair of target allocations. This process can be simplified substantially by accounting for the patrol region of each fibre, which significantly reduces the number of possible combinations from  $\sim 6.9 \times 10^{376}$ .

Compared to the brute force method, the algorithm for Configure is significantly faster; it would otherwise take over a day to compute a single field using an average laptop. The time taken for Configure to assign a single field is dependent on the annealing parameters used and the number of targets in each field; for the default values, this is approximately 8 minutes. Further investigation has been done on the affect of changing the annealing temperatures; however, this is outside the scope of this paper.

The remaining steps for calculating Configure's collision matrix can be divided into four phases. The first phase is to calculate the position of the target in the focal plane, using typical values of atmospheric pressure, temperature, hour angle, and humidity. Each target's location is then used to isolate a list of fibres that can reach it, along with the area that the fibre's button would cover (known as the "footprint"). As this process is independent for each target, Configure parallelizes this calculation for multiple targets simultaneously.

Phase two begins with a series of checks that are independent of the fibre. For example, if two targets are placed so closely that a collision is inevitable, then no further checks are needed for that pair. The checks are then extended to search for overlapping footprints between fibre pairs, and also footprints that overlap other nearby targets. Finally, this phase looks to see if any fibres may overlap the boundary of a button that has been placed on another target. The list of these possible fibre-button crossings are stored so that they may be accessed during the later stages.

The next stage combines the logic of the previous two, by looking at the fibres that can be placed on the first target in a nearby pair of targets. If phase two found that a fibre might overlap the button, it is possible a fibre may overlap the button outline on the second target, then all of the fibres which may reach the first target need to be checked for this condition. The final phase is executed if the previous check returns a conflict. A specific fibre is placed on the first target in the pair, and the same button-fibre conflicts will be checked by changing the fibre placed onto the second target.

By reducing the number of possible fibre target combinations early on in the process, these conflict assessments are significantly faster; when a clash is found, no further tests for conflict are required.

### 3. COMPLETENESS

Understanding how Configure distributes its target assignment across the field of view is scientifically crucial for the WEAVE surveys. Certain areas of each field could be preferentially assigned over others, as a result of either the simulated annealing algorithm, or physical factors such as the fibres' limited range of movement based on their tier. Any regions of low assignment need to be identified, as the completeness would vary as a function of position and this could impact measurements of the luminosity function or space density.

It is important to note that prior to the MOS fibre allocations, each of the eight guide fibres are assigned to specific targets in every field. A minimum of two are required for successful telescope tracking, with one guiding target at the centre of every field that must be assigned. Every survey must also specify how many fibres should be left to observe the night sky with, which are then allocated after the annealing process has finished. For the WL survey, approximately 100 fibres are used as "sky" fibres per field.

Our definition of completeness follows Miszalski et al,<sup>10</sup> who define the completeness,  $C$ , as the fraction of targets assigned over the total number of targets in that region,

$$C = \frac{n_a}{n_a + n_u}, \quad (2)$$

where  $n_a$  is the number of targets allocated, and  $n_u$  is the number of unallocated targets.

Initial studies of Configure<sup>1</sup> recommend that a list of 1400 targets is used per observation to make use of the full fibre multiplex available; this is kept constant for each field in our analysis. As we are not varying the number of targets, we have modified this equation to ignore the additional scale factor in their definition.<sup>10</sup> The field of view was divided into  $5 \times 5$  mm areas (approximately  $89 \times 89''$ ) to ensure that a small number of targets were consistently included in each segment. Reducing these dimensions further leads to segments having one or no targets at all.

We begin by running the 34 fields through configure independently using the default parameters; this is referred to as the nominal case, although here we set the lower priority targets have a value of 5, which will be varied in later sections of this analysis. Figure 5 shows the stacked distribution of all the targets in the 34 mid-tier WL fields used for this study. The colour of every target corresponds to the completeness value computed for the region it lies in, with darker areas of the field indicating consistently low areas of fibre allocation.

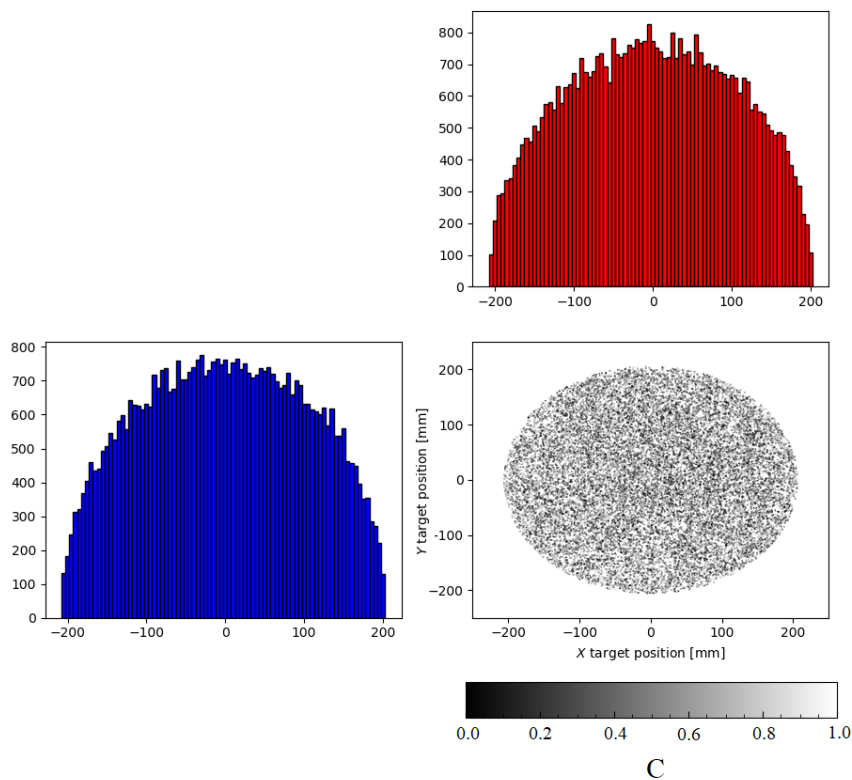


Figure 5. The distribution of targets for all 34 mid-tier WL fields, with 1400 potential targets per field. Each target is shaded by the completeness ( $C$ ) value of the  $5 \times 5$  mm region it falls within. Darker regions of the field indicate areas of low completeness. The blue and red histograms are the distribution targets in the  $x$  and  $y$  direction respectively, for every region.

When first examined, the distribution displayed a faint radial stripe. This implied that fibres in that region had been blocking other targets from being assigned. Further investigation revealed that the Configure algorithm was placing the same guide fibre at the centre of the field each time. As they are allocated prior to the MOS fibre, this meant that a section of the field was permanently inaccessible during the later assignment stages. To rectify this, a guide fibre is now randomly chosen from the list of eight possible fibres.

Having randomised the guide fibre selection, the distribution shows no other discernible features. This is a positive result that is consistent across both plates, meaning that further work is not needed to remove selection effects due to the arrangement of the fibres and the limits of their travel range.

#### 4. TARGET PRIORITY

The nominal assignment case of Configure assigns all targets simultaneously, with an acceptance function that is weighted according to the priority of the target being allocated in each fibre swap. In the annealing algorithm

used by 2dF<sup>10</sup>, this is given by

$$E = \sum_i^{N_{\text{Piv}}} \left[ \beta^{p_i} + \delta \sum_j^{N_{\text{Assoc}}} \beta^{p_j} \right] \left( \frac{\alpha_i - \alpha_{\text{max}}}{\alpha_{\text{max}}} \right)^\gamma, \quad (3)$$

where  $\alpha_{\text{max}}$  is the maximum fibre deviation,  $\alpha_i$  represents the current angle deviation for this pivot,  $N_{\text{piv}}$  is the number of pivots (in this case targets) in the field, and  $p_i$  is the priority of the allocation to the pivot.<sup>10</sup> In this case,  $\beta$  is defined as the weighting base as stated in Campbell et. al 2004.<sup>7</sup> Here,  $\delta$  and  $\gamma$  are real parameters that are used to control the preference for close pair assignment, and favour fibres with small angular deviations (i.e. a radial trajectory).

Configure uses an adapted version of this algorithm, where the energy of each individual fibre with a target assigned to it is given by

$$E = \frac{1 + s}{\text{target priority}}. \quad (4)$$

Here,  $s$  represents the straightness of the fibre, and is much less than 1. The target priority can take any positive integer value, with 0 ensuring that the target is excluded from the assignment process. In all the WEAVE surveys, the maximum value used is 10. This is not limited by Configure itself, and effect of increasing this upper limit is discussed below.

The algorithm then searches for the configuration with the lowest energy. The end temperature of the configuration is an adjustable parameter of Configure. If a fibre swap occurs that lowers the total energy of the system, it will be accepted automatically. However, if the result is an increase in energy, then the swap has a probability  $P$  of being accepted that is defined by

$$P = \exp \left( 1000 \left( \frac{E_2 - E_1}{T} \right) \right), \quad (5)$$

where  $E_2$  is the energy after the swap and  $E_1$  is the energy beforehand. The temperature  $T$  is the current temperature of the system at the time of the swap. As the temperature of the system decreases with time, the probability of swaps being accepted gradually reduces.

The WL science team initially decided to allocate fibres to targets in their fields using only two priority values; a higher priority of 10 for objects crucial to the survey, and a lower priority of 5 for additional targets that they may want to include. In the case of the mid-tier survey, any targets with a flux density above 1 mJy is considered a high priority target. Other WEAVE science teams have taken independent approaches to designating their priority values, which will not be discussed in this paper. For the WL survey, we investigated how the divergence of these priorities impacts the number of higher priority targets allocated per field, and their distribution.

We compare the case where only the lower priority value has been changed from 5 to 1, whilst keeping the highest priority value in each field at 10. Figure 6 shows the number of higher priority targets assigned by Configure against the total number of higher priority targets available in the field, where either a lower priority of 5 or 1 is used for the other targets. This is using the nominal assignment method, where both priorities are allocated simultaneously. We can see from this figure that by reducing the priority value of the less crucial targets to 1, one sees a consistent increase in the proportion of priority 10 targets assigned in each field. Notably, we see that the priority 1 targets are not five times less likely to be observed than the priority 5 targets.

To estimate the degree of variation in the total fibre assignment per field, we ran 450 Monte Carlo<sup>12,14</sup> (MC) realisations for a single mid-tier field, with the lower priority targets assigned a priority value of 1. This number of realisations was chosen to both maximise the accuracy of the uncertainty estimation whilst also reducing the computational time required. These realisations are for the nominal Configure procedure, and the range of outcomes are shown in figure 7.

By taking the uncertainty boundaries as being the 16<sup>th</sup> and 84<sup>th</sup> percentiles of the MC distributions, the uncertainty in the total number of fibres assigned is  $\pm 2$  and for the priority 10 targets only it is  $\pm 9$  fibres. Every

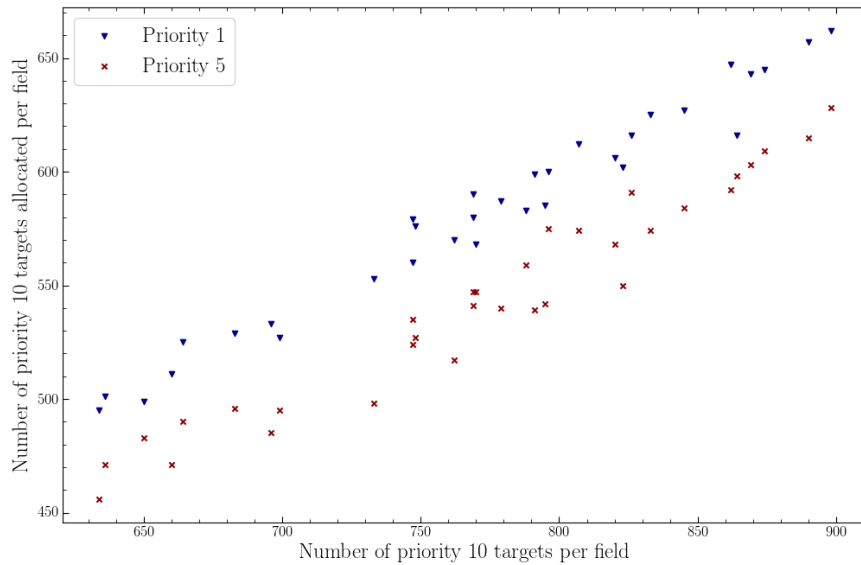


Figure 6. The number of priority 10 targets included in each field against the number allocated by Configure, for each of the 34 fields in the mid-tier WL survey.




Figure 7. Monte Carlo assignment realizations for the nominal Configure usage case, repeating 450 times with a single field. (*left*) The total number of targets assigned for each iteration. (*right*) The number of priority 10 targets assigned per observation, in each realisation. The blue, green, and red lines represent the 16<sup>th</sup>, 50<sup>th</sup>, and 84<sup>th</sup> percentiles respectively.

MOS-mode observation for WEAVE has a number of fibres reserved for sky observations. This is a parameter that each survey can adjust. For our subset of mid-tier WL fields, this is set to 100. The number of sky fibres reserved for each observation will shift the peak of the total number of fibres allocated.

The decision to limit the highest target priority value to 10 was made by the WEAVE science teams to prevent biased assignment in field configurations where different surveys share fibres. As increasing the difference between the two priorities had a notable affect of allocation proportion, it is important to test the limits of this weighting method. The higher priority targets were adjusted to have a value of 10,000, and the lower value was left at




Figure 8. The number of higher priority targets assigned per field, against the number of higher priority targets in each field, for a higher priority value of 10 and 10,000. We see that there is no significant impact on increasing the maximum priority value across the complete set of mid-tier WL fields.

1. From figure 8, one sees that this significant increase in the priority scale does not correspond to a notable increase in the number of high priority targets assigned.

## 5. ASSIGNMENT CASES

The nominal allocation method for each field takes all fibres, regardless of their priority, and begins the annealing process mentioned in section 2.1. It is possible to make adjustments to this process without modifying the algorithm itself. One way is to adjust the annealing parameters, such as the start and end temperatures. However, this may increase the time period Configure must run for. Alternatively, we can modify the input field that Configure is supplied with.

To optimise the performance of Configure, two variations in the assignment process were applied to the mid-tier WL fields by modifying the starting file. We describe the results of implementing a two-stage process, which separates the allocation process by target priority. This method was chosen to prevent the need for additional filler targets in each observation whilst trying to reach the optimal assignment of the higher priority targets. The second test case is to allow all of the targets to be assigned simultaneously, like the nominal case, except that the targets are separated into temporary survey categories by priority. The advantage of this method is that we can limit the number of fibres available to each survey, without having to add additional targets to each field.

### 5.1 Two-stage assignment

The ability to exclude a target from an observation by setting its priority value to zero has many advantages when conducting a survey; for example, someone can exclude targets which are too close together to be configured, or prevent certain targets being assigned multiple times when they appear in overlapping observations.

We use this feature to create a two-stage configuration process. In the first stage, we only allow the priority 10 targets to be assigned in Configure. The output file from this run is then fed through Configure again, containing the fibre allocation information from the first stage and restoring the priority zero values back to their previous non-zero values so that they can be assigned in this second run. Figure 9 shows the results of this process on the number of fibres used per field.

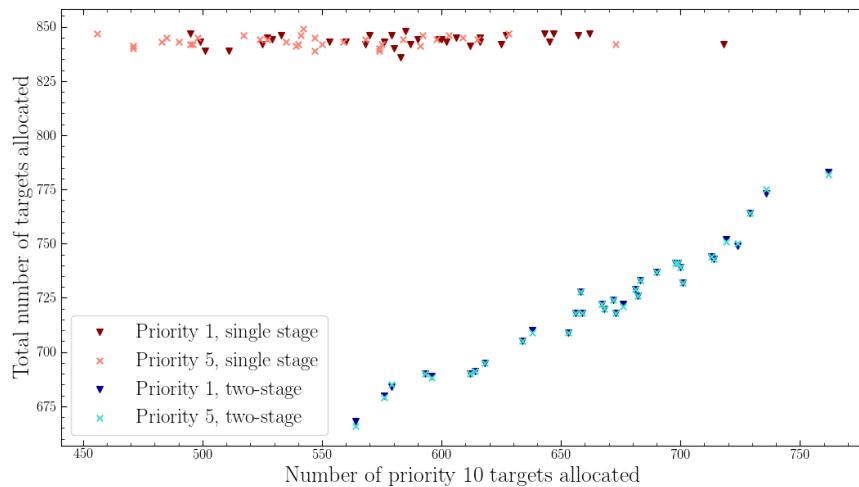


Figure 9. The number of priority 10 targets assigned per field against the total number of targets allocated per field, independent of priority. The cases shown are for a nominal, single stage configuration with either a lower priority value of 1 or 5, or a two-stage assignment process where the lower priority values are either 1 or 5.

This is dissected further to look at the impact of this process on the radial assignment completeness. The set of all fields have been divided into three sub-groups based on the number of priority 10 targets in the field, with boundaries at 700 and 825 priority 10 targets. The completeness as a function of radius is plotted for each of these sub-groups in figure 10.

For the radial completeness, the two-stage method consistently gives a higher assignment fraction as a function of radius. We can see from figure 9 that while using a two-stage process does increase the proportion of higher priority targets assigned, it comes at a cost in terms of the total number of fibres that can be used in each field. Compared to nominal Configure usage, which gives us  $\sim 850$  fibres consistently, the two-stage process gives us only  $\sim 810$  fibres used in the best case. Therefore, to maximise the efficiency of WEAVE, this is not a viable method of target assignment, as the best case results in a decrease in fibre the usage of 5% and 20% in the worst case.

We believe that the reason this method is not as effective is due to a lack of flexibility in the fibre assignment. The first stage "freezes" a large proportion of the fibres into place, which means that a fibre assigned to a higher priority cannot be moved to another high priority target, even if it meant that several lower priority targets could then have fibres allocated to them.

For this reason, we use the first-stage assignment process to provide us with an estimate of the maximum number of higher priority targets that are possible to place in each field. The aim is to then find a method that will bring us as close as possible to this value, without compromising on the number of fibres used per observation.

## 5.2 Multi-survey assignment

Another feature of configure is its ability to assign targets to fields that are shared between multiple science teams. Each target is given a survey name, and each survey has a programmed limit on the number of fibres it is allowed to use for its assignment process.

We wish to test if this feature is more flexible compared to the two-stage configuration, in the hope that we can optimise the output for both the higher priority assignment and the total number of fibres used.

As in the two-stage method, the input file is modified. Two "surveys" are created, one for each priority, which each have a different limit on the number of fibres that can be assigned. For the higher priority survey,

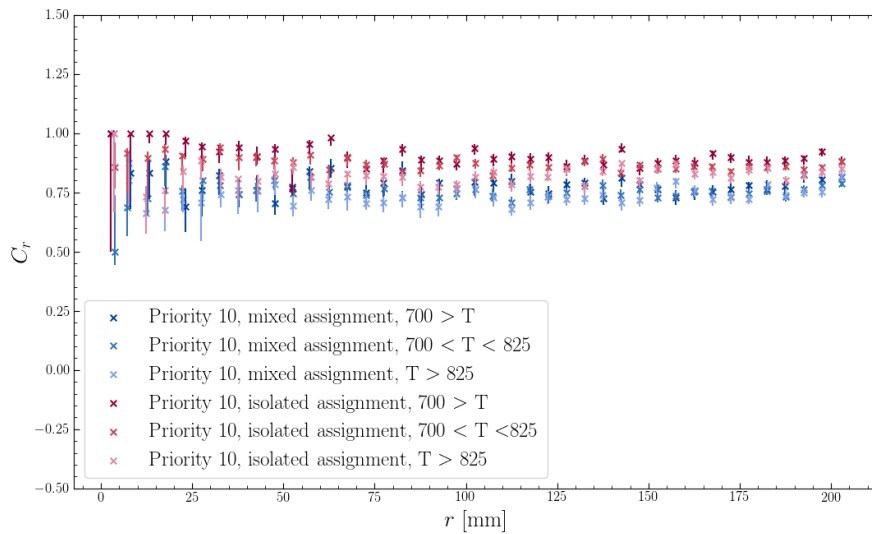


Figure 10. The radial completeness distribution  $C_r$  for both the nominal assignment and the two-stage assignment process. Each case has been divided into three groups based on the number of priority 10 values included in each field, with boundaries at 700 and 825 targets. This distribution shows the average completeness across all fields in each group as a function of radius.




Figure 11. The distribution of assignment for the 34 mid-tier WL fields when configured under the multi-survey priority case. (*left*) The total number of targets assigned in each field, when either 10, 50, or 200 fibres are allowed to be used by the lower priority survey. (*right*) The number of priority 10 targets assigned, when either 10, 50, or 200 fibres are allowed to be used by the lower priority survey.

we set this to be the maximum number of fibres available per field,  $\sim 860$  in total depending on the field plate and the number of sky fibres used. For the lower priority survey, the maximum number of fibres available was varied between 10 to 200 fibres in order to study the effect it had on the allocation statistics.

Figure 11 shows the proportion of high and low priority targets assigned, for the multi-survey method. In this case, each of the 34 mid-tier fields has been processed through configure with a different maximum available for the lower priority survey. We consider the cases where the maximum fibres allowed for the low priority targets are 10, 50, and 200.

Our ideal comparison is the number of higher priority fibres assigned in the first part of the two-stage process. This is considered to be the maximum number of priority 10 targets that can be assigned in the field in the




Figure 12. Monte Carlo assignment realisations for the multi-survey Configure usage case, repeated 450 times with a single field where 200 fibres are allowed for the lower priority survey. (*left*) The total number of targets assigned for each iteration. (*right*) The number of priority 10 targets assigned for each iteration. The blue, green, and red lines represent the 16<sup>th</sup>, 50<sup>th</sup>, and 84<sup>th</sup> percentiles respectively.

absence of any other targets. Although, in practise Configure does not find the global maximum, which leads to a degree of variation. Whilst reducing the number of fibres available for the lower survey does shift us closer towards this ideal value, it still reduces the total fibres used per observation. The difference between the peak of the ideal targets allocated for all fields and the maximum fibres available is approximately 200. We use this value to limit the lower survey, and see if this produces an improvement on the number of higher priority assignment compared to the nominal case without impacting the total fibres used.

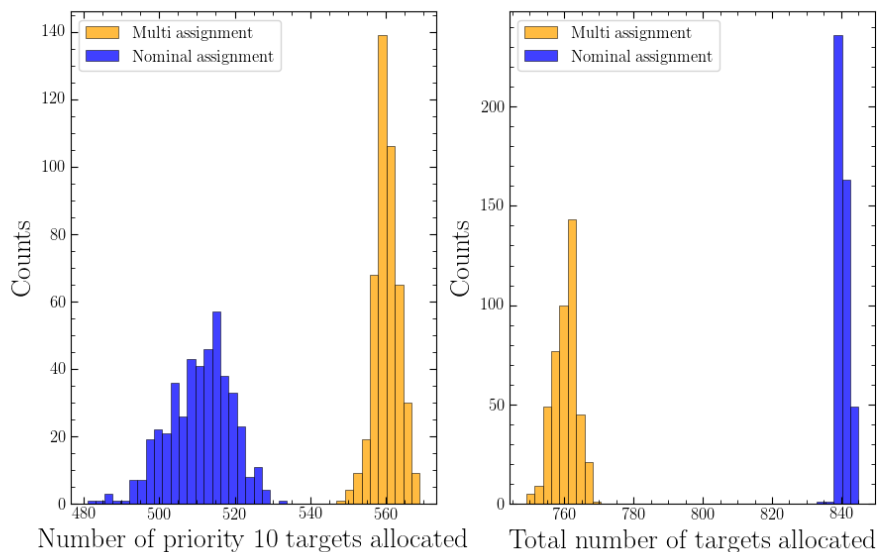


Figure 13. MC realisations of the same field, processed by Configure 450 times for both the nominal case and using the multi-survey method. (*left*) Shows the distribution of priority 10 targets allocated per realisation. (*right*) Displays the total number of targets allocated in each realisation.

Monte-Carlo simulations were used to observe the variation of the assignment process. The same field was configured using the multi-survey method for 450 realisations, keeping the maximum fibre limits constant. The

results are shown in figure 12, where the 16<sup>th</sup> and 84<sup>th</sup> percentiles are taken as the upper and lower error bounds. These were found to be  $\pm 3$  for both the total and higher target allocation.

In figure 13, we overlay the impact of the multi-survey method with the same field in the nominal case, that were repeated for 450 realisations. It should be noted that this is a different field compared to the realisations in figure 12, which has a distinct configuration setup and number of higher priority targets. For the total number of targets assigned, the multi-survey case has a median which is 80 fibres lower compared to the nominal case. However, when considering the higher priority targets only, the assignment distribution shows that the multi-survey process has a peak which is shifted by 49 targets compared to the nominal case. Whilst this does result in an average decrease of 9.5% in the total fibres used from the nominal assignment, there is an average increase of 9.5% in the number of higher priority targets assigned, and a 12.3% increase in the best case scenario. Notably, this method does not use any additional filler targets to achieve this affect. As a result, this method should be considered as an alternative starting point for further analysis, which appears promising compared to the nominal assignment case.

## 6. FLUX RESTRICTIONS

It is important to understand Configure's assignment process in the context of the flux density distribution for the WL catalogue. This has not been previously analysed by any of the WEAVE science teams. Doing so may highlight unknown selection effects, which will impact the scientific analysis conducted using WEAVE data e.g. stellar and galactic luminosity functions. For this study we use the subset of mid-tier WL fields with a nominal assignment procedure, where targets with a flux density  $> 1$  mJy are given a priority value of 10, and targets with a flux density  $< 1$  mJy are given a value of 1. The results are dependent on the survey fields used, however the overall behaviour is relevant for all MOS observations across the WEAVE surveys, when the number density of potential targets correlates with decreasing flux (or magnitude).




Figure 14. The cumulative fraction of targets assigned above the flux threshold value, compared to the flux value for each target. This is normalised by the total number of targets available in the field, and by the complete catalogue distribution of targets (represented by the black line). Each distribution has a selection cut ranging from 1.0 mJy to 0.6 mJy.

We begin by applying a series of selection cuts to each of the 34 fields in the mid-tier subset, based on the flux density, ranging from 1.0 to 0.6 mJy. Targets above this threshold are included in the assignment process and do not have their priority value changed. Therefore, both high and low priority targets are included in the assignment if they have a high enough flux density. For this analysis, the lower priority targets have a value of 1 if they are included. Any target with a flux below this threshold is excluded by setting its priority value to

zero. After being processed, every field is run through `configure` once for each selection cut. The targets are then cross-matched with their flux value from the catalogue.

We wish to analyse the completeness of the assignment process as a fraction of the targets available in the catalogue, to determine whether it is flux dependent. The completeness in this case is defined as

$$\left( \frac{N(F_i \geq f)}{N_{\text{available}}} \right) / N_{\text{catalogue}}, \quad (6)$$

where  $N(F_i \geq f)$  is the reverse cumulative sum of the assigned targets as a function of their flux density. For this analysis,  $N_{\text{available}}$  is kept constant at 850 fibres for each of the 34 fields, and  $N_{\text{catalogue}}$  is the number of targets in the WL mid-tier catalogue. This is shown in figure 14. A logarithmic scale is used for the flux density to display the full catalogue range.

Decreasing the selection cut threshold means that there is a larger combination of both high and low priority targets included in the assignment process of each field. In figure 14, the selection cut of 0.6 mJy demonstrates a gradual downward curve, that indicates a flux dependent completeness. Conversely, the selection cut of 1.0 mJy, displays a curved upward trend as a function of flux density. This behaviour is counter-intuitive to the behaviour we would expect, and demonstrates that the flux dependence varies with the configuration arrangement. This selection effect must be accounted for in future observational analysis.

At the brighter end of the flux density distributions, the Poisson errors dominate due to a reduction in the number of targets on this scale. We also observe sharp variations in the completeness as a result of `Configure` finding the local minimum in the annealing procedure.

## 7. OVERLAPPING FIELDS

To efficiently tile the whole northern sky, WL<sup>2</sup> will use a pointing strategy based on that of Saff & Kuijlaars.<sup>15</sup> Using this strategy, approximately 88% of targets will fall within one WEAVE field of view. The remaining targets will either be visited twice or not at all.

Total number of targets in this subset	47600	%
Targets included in a single field	44891	94.31
Targets included in multiple fields	2709	5.69

Table 1. The number of targets which are included in the allocation process, across all 34 fields, as well as the proportion of these targets that appear in one or more fields.

Targets included in one field	44891	%
Unassigned	19881	44.29
Assigned once	25010	55.71
Targets included in two fields	2709	%
Unassigned	329	12.14
Assigned once	1088	40.16
Assigned twice	1292	47.70

Table 2. The number of targets included in single or multiple fields for the nominal assignment case, and their respective allocation values.

In this case, an observation is defined as the field of view that a target may fall within, based on the tiling pattern for this subset of fields. For the purposes of this study, the expression *included* means that a target is allowed to be allocated in `Configure`'s processing of a field. Additionally, the expression *assigned* means that a target that is included in the field allocation process has successfully had a fibre placed on its location. As a result, a target may fall within an observation, and still not be included in the field allocation process; or a target may be included in a field, and it hasn't been allocated a fibre.

The assignment, as opposed to the visit statistics, for the overlapping regions of each field need to be simulated, to understand the impact of `Configure` on the target coverage across the survey. This includes calculating the

proportion of targets that are included in one or more Field of View for this subset of fields, as shown in table 1, and the number of times a target is assigned, particularly if it is included in more than one field (tables 2 & 3).

For all targets in the mid-tier WL survey catalogue, the assignment statistics are summarised in table 2 and presented in figure 15 for the nominal case. As expected, the majority of the targets included in a single field of view are assigned, as 840 fibres are assigned for each realisation on average.

Interestingly, for targets included in two fields, it seems that they are almost equally likely to either be assigned once as they are twice, with a significantly smaller proportion of targets remaining unassigned. This is not unexpected, as these targets are located towards the edge of the field of view. If a target is easily accessible in one field, then it is likely to also be accessible in the second field of view.




Figure 15. The proportion of the mid-tier WL catalogue targets that are either not assigned a fibre (“excluded”), or are assigned a fibre “included” in one or more fields. Alongside this is the proportion of the targets that are assigned once, twice, or remain unassigned by Configure. This is for the nominal case only.

Targets included in one field	44891	%
Unassigned	20845	46.43
Assigned once	24046	53.57
Targets included in two fields	2709	%
Unassigned	450	16.61
Assigned once	944	34.85
Assigned twice	1315	48.54

Table 3. The number of targets included in single or multiple fields for the multi-survey case, and their respective allocation values.

We expand this further to compare the differences between the nominal and multi-survey methods. The multi-survey assignment proportions listed in table 3. For targets included in a single field of view, we find that the assignment percentages only vary by  $\sim 1\%$  from the nominal method. Where targets are included in two fields, there is a  $\sim 4\%$  increase in the number of targets unassigned in the multi-survey case, however there is an increase of  $\sim 1\%$  in the number of fibres assigned twice. This is at the expense of fibres which are only assigned once, which implies that there is more flexibility in the assignment of the overlapping regions for the multi-survey case. This comparison is shown in further detail in figure 16 for all 34 fields for each priority value used. This

figure indicates that the multi-survey method assigns a larger number of higher priority targets twice compared to the nominal method, that is reflected in the proportion of lower priority targets allocated.

Further MC realisations of two overlapping fields were run for the multi-field case, shown in figure 17. This is isolated to the overlapping field region, to show the variation in the number of priority 10 targets allocated. We can clearly see that this distribution is not consistent between the two fields. This implies that targets located towards the centre of the field affect the assignment process at the edges. As a result, the allocation of the shared regions is highly subjective to each survey’s configuration setup.




Figure 16. The distribution of target assignments across the 34 mid-tier WL survey fields used for this study. This is separated by priority for both the nominal assignment, and the multi-survey assignment cases. Only targets that are included in at least one observational field are shown, however this work can be expanded to cover the full survey.

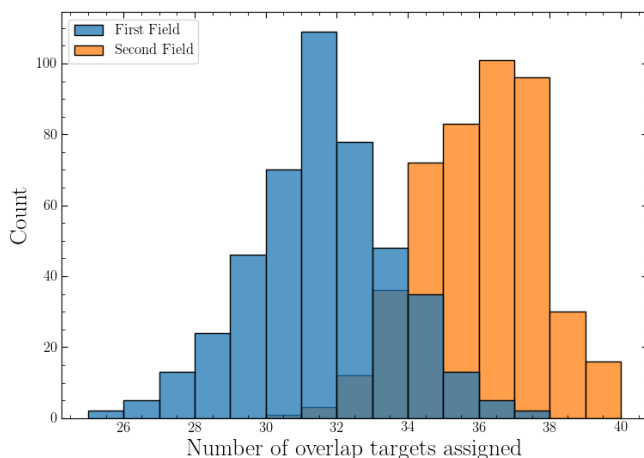


Figure 17. MC realisations for two observations that share a region of the field. We show the variation of the number of priority 10 targets for only the overlapping area of the field. We see that the distribution is not consistent between fields of view.

## 8. CONCLUSION AND FUTURE WORK

From this work, it is clear that Configure uses a very complex method of fibre assignment that cannot be controlled in detail when optimising the assignment of WEAVE's  $\sim 960$  fibre multiplex to the 1400 targets contained in each proto-field. We have explored different methods of using Configure, and whilst they do not provide a complete solution, they do provide an additional element of control that can be utilised by the WEAVE survey teams.

Features of Configure, such as the priority weighting and the simultaneous survey assignment process, can be used to improve the fibre allocation output. We find that changing the difference between priorities of two target types can improve the total fraction of assigned fibres per field, however scaling this difference beyond the value of 10 did not show any significant improvements.

By using a two-stage assignment method, where the field is processed twice, we do see an increase in the proportion of higher priority targets that are allocated fibres. This comes at a significant cost to the total number of fibres used per field. Therefore, is not an acceptable method of assignment, as the observational capabilities of the WEAVE fibre positioner would not be fully utilised.

An alternative method is to use Configure's capability to simultaneously assign targets that are shared between survey's. We modify the original field file to separate the targets into two surveys based on their priority value. For the higher priority survey, we set the number of fibres available to the maximum that can be used in a single field. For the lower priority survey, the optimal number of fibres with this subset of WL fields is 200, although this will vary for alternative configuration strategies. These 200 fibres are the difference between the "ideal" number of higher priority fibres that we believe can be allocated in a single field, and the total number of fibres available in each configuration. Whilst reducing the number of fibres available for the lower priority targets does increase the proportion of higher targets selected, this still leaves a notable amount of fibres that are left unused.

Monte Carlo simulations were completed with 450 realisations through Configure for a single field in both the nominal and multi-survey arrangement. This compromise on the number of fibres allowed in the lower priority survey still results in an overall increase in the number of higher priority allocated, characterised by a shift in the peak of the allocated distribution from the nominal case.

The assignment values across field overlap regions were compared against the complete target selection catalogue. The majority of all targets were included in at least one survey field, however there is a greater quantity of fibres that are included in one field and remain unassigned than there are targets which do not fall within any field of view. We also find that the targets included in an overlapping region are almost as likely to be assigned once as they are to be assigned twice, over the mid-tier survey.

With our new priority scheme, we studied impact of the flux selection on the targets included and assigned in each observation. Filtering targets for a flux threshold that ranged from 1.0 to 0.6 mJy in the WL catalogue resulted in a cumulative distribution that separates on the lower flux scale. We observe a gradual curve, in opposing directions depending on the selection cut, as the targets flux increases that implies a dependence on the flux density. This is an important result that must be accounted for when conducting further analysis using WEAVE data, such as stellar and galactic luminosity functions. The reduced number of higher flux targets in the survey causes an increase in the associated Poisson errors and stronger variations in the assignment distribution.

In future, this work should be expanded to other WEAVE survey fields to assess the impact of Configure on their predicted science output and target inclusion. When the final version of Configure is released, the allocation distribution should be checked to ensure that the presence of disabled fibres does not introduce regions of low completeness into the observation.

Further experiments can be done by comparing Configure with more recent methods of solving combinatorial optimisation problems, as well as by adjusting the annealing parameters themselves. Genetic algorithms and reinforcement learning processes have become very efficient at finding alternative heuristic solutions. Further information on them can be found in Grefenstette et al.<sup>16</sup> and Gambardella et al.<sup>17</sup>

## ACKNOWLEDGMENTS

This work was supported by the Science and Technology Facilities Council and St. Cross College, Oxford. The author would like to give special acknowledgement to SPIE for their support in funding their attendance to this conference.

Funding for the WEAVE facility has been provided by UKRI STFC, the University of Oxford, NOVA, NWO, Instituto de Astrofísica de Canarias (IAC), the Isaac Newton Group partners (STFC, NWO, and Spain, led by the IAC), INAF, CNRS-INSU, the Observatoire de Paris, Région Île-de-France, CONCYT through INAOE, Konkoly Observatory of the Hungarian Academy of Sciences, Max-Planck-Institut für Astronomie (MPIA Heidelberg), Lund University, the Leibniz Institute for Astrophysics Potsdam (AIP), the Swedish Research Council, the European Commission, and the University of Pennsylvania. The WEAVE Survey Consortium consists of the ING, its three partners, represented by UKRI STFC, NWO, and the IAC, NOVA, INAF, GEPI, INAOE, and individual WEAVE Participants. The WEAVE website can be found at <https://ingconfluence.ing.iac.es/confluence//display/WEAV/The+WEAVE+Project> and the full list of granting agencies and grants supporting WEAVE can be found at <https://ingconfluence.ing.iac.es/confluence/display/WEAV/WEAVE+Acknowledgements>.

## REFERENCES

- [1] Terrett, D., “Fibre positioning algorithms for the weave spectrograph,” in [*Software and Cyberinfrastructure for Astronomy III*], *Proc. SPIE* **9152** (2014).
- [2] Smith, D., Best, P., Duncan, K., Hatch, N., Jarvis, M., Röttgering, H., Simpson, C., Stott, J., Cochrane, R., Coppin, K., et al., “The weave-lofar survey,” in [*SF2A-2016: Proceedings of the Annual meeting of the French Society of Astronomy and Astrophysics*], **1** (2016).
- [3] Dalton, G., “Final design and progress of weave: the next generation wide-field spectroscopy facility for the william herschel telescope,” in [*Ground-based and Airborne Instrumentation for Astronomy VI*], *Proc. SPIE* **9908** (2016).
- [4] Dalton, G., “Integration and early testing of weave: the next-generation spectroscopy facility for the william herschel telescope,” in [*Ground-based and Airborne Instrumentation for Astronomy VIII*], *Proc. SPIE* **11447** (2020).
- [5] Hughes, S., “Final assembly, metrology, and testing of the weave fibre positioner,” in [*Ground-based and Airborne Instrumentation for Astronomy VIII*], *Proc. SPIE* **11447** (2020).
- [6] Lewis, I. J., Cannon, R. D., Taylor, K., Glazebrook, K., Bailey, J. A., Baldry, I. K., Barton, J. R., Bridges, T. J., Dalton, G. B., Farrell, T. J., Gray, P. M., Lankshear, A., McCowage, C., Parry, I. R., Sharples, R. M., Shortridge, K., Smith, G. A., Stevenson, J., Straede, J. O., Waller, L. G., Whittard, J. D., Wilcox, J. K., and Willis, K. C., “The Anglo-Australian Observatory 2dF facility,” *Monthly Notices of the Royal Astronomical Society* **333**, 279–299 (June 2002).
- [7] Jin, S. et al., “The weave survey: Design, overview and simulated implementation,” *Monthly Notices of the Royal Astronomical Society in preparation* (2022).
- [8] Kirkpatrick, S., C. D. Gelatt, J., , and Vecchi, M. P., “Optimization by simulated annealing,” *Science* **220**, 671–680 (1983).
- [9] van Haarlem, M. P., Wise, M. W., Gunst, A., Heald, G., McKean, J. P., Hessels, J. W., de Bruyn, A. G., Nijboer, R., Swinbank, J., Fallows, R., et al., “Lofar: The low-frequency array,” *Astronomy & astrophysics* **556**, A2 (2013).
- [10] Miszalski, B., Shortridge, K., Saunders, Parker, Q. A., and Croom, S. M., “Multi-object spectroscopy field configuration by simulated annealing,” *Monthly Notices of the Royal Astronomical Society* **371**, 1537–1549 (2006).
- [11] Schrijver, A., [*Combinatorial Optimization: Polyhedra and Efficiency*], Springer Science & Business Media (2003).
- [12] Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., and Teller, E., “Equations of state calculations by fast computing machine,” *J. Chem. Phys.* **21**, 1087–1091 (1953).
- [13] Berger, J. O., [*Statistical decision theory and Bayesian analysis*], Springer (1985).

- [14] Robert, C. and Casella, G., “A short history of markov chain monte carlo: Subjective recollections from incomplete data,” *Statistical Science - STAT SCI* **26** (08 2008).
- [15] Saff, B. E., Kuijlaars, and B. J. A., “Distributing many points on a sphere,” *The mathematical intelligencer* **19**(1), 5–11 (1997).
- [16] Grefenstette, J., Gopal, R., Rosmaita, B., and Van Gucht, D., “Genetic algorithms for the traveling salesman problem,” in [*Proceedings of the first International Conference on Genetic Algorithms and their Applications*], **160**(168), 160–168, Lawrence Erlbaum (1985).
- [17] Gambardella, L. M. and Dorigo, M., “Ant-q: A reinforcement learning approach to the traveling salesman problem,” in [*Machine learning proceedings 1995*], 252–260, Elsevier (1995).