



RESEARCH ARTICLE

10.1029/2023SW003487

SIR-HUXt—A Particle Filter Data Assimilation Scheme for CME Time-Elongation Profiles

Luke Barnard¹ , Mathew Owens¹ , Chris Scott¹ , Matthew Lang¹ , and Mike Lockwood¹ 

¹University of Reading, Reading, UK

Key Points:

- A sequential importance resampling (SIR) data assimilation scheme is coupled to the HUXt solar wind model, called SIR-HUXt
- Observing system simulation experiments are used to study the impact of assimilating heliospheric imager Coronal Mass Ejection (CME) observations into SIR-HUXt
- These experiments show SIR-HUXt can significantly reduce the uncertainty on CME arrival time and arrival speed simulations

Correspondence to:

L. Barnard,
l.a.barnard@reading.ac.uk

Citation:

Barnard, L., Owens, M., Scott, C., Lang, M., & Lockwood, M. (2023). SIR-HUXt—A particle filter data assimilation scheme for CME time-elongation profiles. *Space Weather*, 21, e2023SW003487. <https://doi.org/10.1029/2023SW003487>

Received 16 MAR 2023

Accepted 22 MAY 2023

Author Contributions:

Conceptualization: Luke Barnard
Formal analysis: Luke Barnard
Funding acquisition: Mathew Owens, Chris Scott
Methodology: Luke Barnard, Matthew Lang
Project Administration: Mathew Owens, Chris Scott
Software: Luke Barnard, Mathew Owens
Supervision: Mathew Owens, Chris Scott, Mike Lockwood
Writing – original draft: Luke Barnard
Writing – review & editing: Luke Barnard, Mathew Owens, Chris Scott, Matthew Lang, Mike Lockwood

© 2023. The Authors.

This is an open access article under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

Abstract We present SIR-HUXt, the integration of a sequential importance resampling data assimilation scheme with the HUXt solar wind model. SIR-HUXt assimilates the time-elongation profiles of Coronal Mass Ejection (CME) fronts in the low heliosphere, like those extracted from heliospheric imager (HI) data. Observing System Simulation Experiments are used to explore SIR-HUXt's performance for a simple synthetic CME scenario of an Earth directed CME in a uniform solar wind, where the CME is initialized with the average CME speed and width. These experiments are performed for a range of observer locations, from 20° to 90° behind Earth, spanning the L5 point where ESA's Vigil mission will return HI data for operational space weather forecasting. For this idealized scenario, SIR-HUXt performs well at constraining the CME speed, and has some success at constraining the CME longitude while the CME width is largely unconstrained by SIR-HUXt. Rank-histograms suggest the SIR-HUXt ensembles are well calibrated, with no indications of bias or under/over dispersion. Improved constraints on the initial CME speed lead to improvements in the CME transit time and arrival speed. For an L5 observer, SIR-HUXt reduced the transit time and arrival speed uncertainties by 69% and 63%. Therefore, SIR-HUXt could improve the real-world representivity of HUXt simulations and reduce the uncertainty of CME arrival time forecasts. The idealized scenario studied here likely enhances SIR-HUXt's performance relative to the challenge of simulating real-world CMEs and solar wind conditions. Future work should validate SIR-HUXt with case studies of real CMEs in structured solar wind.

Plain Language Summary Coronal Mass Ejections (CMEs) are eruptions of plasma from the Sun and are the main driver of space weather at Earth. We need to forecast the arrival of CMEs at Earth, to enable actions to mitigate space weather impacts. There is significant uncertainty in CME forecasts because it is difficult to estimate the CME properties needed for the forecast simulations, which limits the usefulness of CME forecasts. In this work we describe SIR-HUXt, a coupled data assimilation scheme and solar wind model. SIR-HUXt's objective is to reduce the uncertainty in CME forecasts by comparing them with estimates of the CMEs location derived from spacecraft images of the Sun. We present a set of theoretical experiments that show that SIR-HUXt can reduce the uncertainty on the CME arrival time and speed estimates, dropping by up to 69% and 63% for an observer at the L5 Lagrange point, where ESA's Vigil space weather monitor mission will operate. The idealized scenario studied here likely enhances SIR-HUXt's performance relative to the challenge of simulating real-world CMEs and solar wind conditions. Future work should validate SIR-HUXt with case studies of real CMEs in structured solar wind.

1. Introduction

Coronal Mass Ejections (CMEs) are large eruptions of magnetized plasma from the Sun's atmosphere. CMEs are the primary cause of severe and extreme space weather at Earth, and so understanding the heliospheric evolution of CMEs, and forecasting their arrival at Earth receives a significant amount of research effort. Currently, numerical simulations of the real-world heliospheric evolution of CMEs have significant uncertainty. This is partly evidenced by the large average errors in CME arrival time of ± 10 hr (Riley et al., 2018). These uncertainties are caused by a range of factors, but appear to be dominated by uncertainties in the initial and boundary conditions of the solar wind models, with uncertainty in the ambient solar wind structure and CME parameters introducing similar levels of uncertainty (Riley & Ben-Nun, 2021). These uncertainties therefore limit our ability to draw scientific conclusions from the simulations of real-world CMEs, and limit the skill of space weather forecasts of CMEs.

This motivates the pursuit of methods with which to reduce the uncertainty on numerical simulations of real-world CMEs. Data Assimilation (DA) methods have excellent potential for improving the representivity of solar wind

numerical models. The objective of DA is to combine information from simulations and observations to provide an optimal estimate of the state of a dynamical system. Heliospheric DA is still a relatively new research topic, but progress is beginning to be made.

Lang et al. (2017) explored how the Local Ensemble Transform Kalman filter could be used to assimilate in situ observations of solar wind plasma properties into the ENLIL magnetohydrodynamic (MHD) solar wind model, which demonstrated clear improvements in the representivity of the ENLIL simulations. The Burger Radial Variational Data Assimilation (BRaVDA) scheme was developed in Lang and Owens (2019), in which a variational DA scheme was coupled to the hydrodynamic (HD) HUX solar wind model (Riley & Lionello, 2011), for the assimilation of observations of the solar wind speed. Experiments with synthetic observations and solar wind speed observations from the STEREO spacecraft showed that BRaVDA reduced the errors in the solar wind speed predictions at Earth. This work was extended by Lang et al. (2021) to the HUXt model, a HD solar wind model with explicit time-dependence (Barnard & Owens, 2022; M. J. Owens et al., 2020), in which it was shown that over the period 2007–2014, BRaVDA returned a 31% reduction of the root mean square error in hindcasts of the solar wind speed at Earth.

These works have so far focused on the assimilation of in situ observations of solar wind plasma properties, but progress has also been made on the assimilation of remote sensing observations, such as those provided by HIs (Eyles et al., 2008; Howard et al., 2008) and interplanetary scintillation (IPS) (Fallows et al., 2022). For example, Barnard et al. (2020) showed that an ensemble of solar-wind-CME simulations with the HUXt model could be weighted by the time-elongation profiles of CMEs derived from the STEREO HI data. This weighting prioritized ensemble members that more closely matched the observed time-elongation profile, and led to up to 20% improvements in hindcasts of the CMEs arrival time at Earth. Similarly, Iwai et al. (2021) demonstrated how assimilating IPS observations of 12 halo CMEs into the SUSANOO-CME MHD model led to improvements in the predicted Earth arrival times of these CMEs.

Although Barnard et al. (2020) demonstrated that HI data contains useful information on CMEs that can be used to constrain the HUXt solar wind simulations, they did not use formal DA methods. In this work, we present the development of SIR-HUXt, which couples a sequential importance resampling (SIR) particle filter DA scheme with the HUXt solar wind model. SIR-HUXt is constructed to assimilate time-elongation profiles of a CMEs flank, such as those typically extracted from the STEREO-HI data (Barnard et al., 2015, 2017; Davies et al., 2009). This is an important milestone toward the development of DA schemes that can directly assimilate the HI intensity data into solar wind numerical models. We present a first test of SIR-HUXt by using Observing System Simulation Experiments (OSSEs) to investigate the performance of SIR-HUXt for a simple synthetic CME scenario and for a range of observer locations relative to Earth.

This article proceeds with Section 2 describing the models and methods we use, including the HUXt numerical model, the background to the SIR algorithm, and on OSSEs. Section 3 presents the results of the OSSEs, and our conclusions are presented in Section 4.

2. Methods and Data

2.1. HUXt

HUXt is an open source numerical model of the solar wind, developed in Python (Barnard & Owens, 2022; M. J. Owens et al., 2020). It is a 1D radial model that uses a reduced-physics approach to produce solar wind simulations that emulate the solar wind flows produced by 3-D MHD models, but at a small fraction of the computational cost.

The motivation for developing HUXt is that the models simplicity and computational expense permits the development of certain experiments and techniques that would typically be too expensive with 3-D MHD models. For example, the particle filter DA experiments in this study require $\approx 10^6$ 5-day simulations of the inner heliosphere, which is currently an impractical demand of 3-D MHD solar wind models with widely available computing resources.

Being based on incompressible hydrodynamics, HUXt solves only for the solar wind flow speed. Consequently, the only boundary condition required is the flow speed at the inner boundary. These boundary conditions can be computed from a wide range of coronal models, including but not limited to; potential field source surface based

models, such as WSA (Arge & Pizzo, 2000) and DUMFRIC (Yeates et al., 2010); MHD models such as MAS (Riley et al., 2001); and tomographically derived conditions, such as CORTOM (Bunting & Morgan, 2022). The combined efficiency and flexibility of HUXt has enabled its development and integration into an upcoming operational ensemble solar wind modeling framework for the UK Met Office, named the Space Weather Empirical Ensembles Package. This provides additional motivation for developing the SIR HUXt DA scheme.

In this work, HUXt is run with its default configuration. The radial grid spans $30 R_{\odot}$ to $240 R_{\odot}$, with a grid step of $1.5 R_{\odot}$. The time-step is 8.7 min. There are 128 evenly spaced longitudinal bins, although to save on computation, and as we are only examining Earth-directed CMEs, the simulation domain only spans the longitude range of $\pm 70^{\circ}$.

CMEs are included in HUXt via the Cone CME parameterization, in which CMEs are represented as a time-dependent velocity perturbation to the model inner boundary. Six parameters are required to specify the initiation of a Cone CME; the initiation time; the speed; the angular width; the source longitude and latitude; and the radial thickness of the perturbation. Further details of the Cone CME parameterization in HUXt are given in M. J. Owens et al. (2020), Barnard et al. (2021), Barnard and Owens (2022). CMEs are tracked through HUXt simulations by inserting test particles into the flow on the CME surface at the model inner boundary. These test particles then passively advect with the flow and are followed at all time steps out to the model's outer boundary.

Pseudo-observers are used with HUXt to compute the time-elongation profile of the Cone CME flank, to emulate the time-elongation data products typically derived from HI observations for example, (Barnard et al., 2015, 2017; Davies et al., 2009; Pant et al., 2016). This is achieved by computing the elongation of each particle on the CME boundary and finding the particle with maximum elongation in an observer's field of view. A better solution would be to forward model the observations from Heliospheric Imager instruments by performing Thomson scattering simulations with HUXt output. However, the HUXt equations are derived from incompressible hydrodynamics, and so only the flow speed is solved for, not the flow mass density. This prohibits a fully self-consistent forward modeling of HI data from HUXt simulations. Consequently, tracking the maximum elongation of the CME tracer particles is a necessary approximation. However, both Barnard et al. (2020) and Chi et al. (2021) showed that this approach returned time-elongation profiles that compared favorably to those extracted directly from STEREO-HI images, which gives us confidence that this approximation is reasonable.

2.2. Sequential Importance Resampling (SIR)

The objective of DA is to provide an optimal estimate of the state of a system by combining the information from both a model and observations of the system, taking proper account of the uncertainties on each.

This can be expressed mathematically via Bayes' theorem, which states that,

$$p(\psi|\theta) = \frac{p(\theta|\psi)p(\psi)}{p(\theta)}. \quad (1)$$

The factors in this equation are typically separated into several colloquially named terms. The “prior,” $p(\psi)$, is the probability density of the model being in a specific state, in the absence of any other external information. The “likelihood,” $p(\theta|\psi)$, which is the probability density of obtaining a set of observations θ , given a model state ψ . The “evidence,” $p(\theta)$, is the probability density of obtaining a set of observations although in most practical examples the evidence becomes a normalizing constant that can be ignored. Finally, the “posterior,” $p(\psi|\theta)$, is the conditional distribution of model states given a set of observations.

Computation of the posterior, or approximations to it, is the focus of DA. The posterior provides the optimal estimate of the state of the system, representing the distribution of model states that are most consistent with the observations. In practical geophysical examples, it is not possible to fully characterize the posterior distribution, and different DA methodologies are used to infer certain properties of the posterior for example, its mean, mode, or variance (Burgers et al., 1998; Le Dimet & Talagrand, 1986). Particle filters are a set of a DA methodologies that aim to approximate the full posterior distribution via an ensemble of “particles” (Ades & van Leeuwen, 2013; Browne & van Leeuwen, 2015; Chorin & Tu, 2009; Fearnhead & Künsch, 2017; Potthast et al., 2019; Van Leeuwen, 2009). For the avoidance of any confusion, the “particles” discussed in the context of DA and SIR are unrelated to the tracking particles used in HUXt, but there is an unfortunate overlap in the nomenclature of these different topics.

SIR is a method of particle filtering that can be used for sequential DA (Fearnhead & Künsch, 2017; Van Leeuwen, 2009). In SIR, the posterior is approximated by the analysis of an ensemble of simulations, or “particles.” The prior is approximated by generating an ensemble of simulations that reflects the uncertainty in the models initial and boundary conditions. The model evolves the ensemble forward in time, until a set of observations are available. At the observation time, an analysis is performed which weights each simulation in accordance with its agreement with the observations. Then, this weighted ensemble is used to generate a new ensemble of simulations which are closer to the observations. The model then resumes advancing the simulations forward in time, until the next set of observations are available. The DA proceeds in this way, performing sequential analysis steps when observations are available. The posterior distribution, at some specific time, is approximated by the distribution of the ensemble after an analysis step.

In this work we develop SIR-HUXt, a coupling of an SIR scheme with the HUXt solar wind model, with the objective of assimilating the time-elongation profiles of CME fronts, such as those that can be derived from white light heliospheric imaging. SIR-HUXt essentially functions as a form of parameter estimation, returning estimates of the posterior of the Cone CME parameters that are most consistent with the observed time-elongation profile of a CME. The following subsections describe the specifics of the SIR algorithm used in SIR-HUXt.

2.2.1. Initial Ensemble Generation

The initial ensemble is generated by perturbing a subset of the Cone CME parameters only, following a similar method to Barnard et al. (2020). Specifically, perturbations are applied to the Cone CME speed, angular width, and longitude. We focus on these three parameters only as they are probably both the most important and most uncertain parameters for determining if and when a CME impacts Earth (Pizzo et al., 2015; Riley & Ben-Nun, 2021), whilst considering all the Cone CME parameters would be too computationally expensive for this proof-of-concept study. Conversely, the CME initiation time is much less uncertain as it is better constrained by the coronagraph data, and the CME thickness is less important as it has a more marginal impact on the CME dynamics in HUXt.

The random perturbations for each parameter are drawn from a uniform distribution that represents the observational uncertainty on that parameter, and the perturbation is added to the best-guess of the true Cone CME parameter. For the speed, width, and longitude, the spread of the uniform perturbation distributions is $\pm 10\%$, $\pm 5^\circ$, and $\pm 5^\circ$, respectively.

The true uncertainty distributions are unlikely to be uniform, but there is not yet good knowledge on what form the observational uncertainties take. And so, in the absence of better knowledge, we follow Barnard et al. (2020) and use the uniform distribution.

The size of the ensemble must be large enough to avoid “filter degeneracy,” where the ensemble essentially collapses into one particle. This occurs when one particle has much larger weight than others in the ensemble such that it dominates the resampling procedure, leading to new particles that are degenerate. The upper bound on the sample size is determined by the availability of computational resources. Here we use an ensemble size of 50. This was determined empirically during our initial experiments. Future work should look to optimize the ensemble size, but for our purposes 50 members perform sufficiently well.

2.2.2. Particle Weighting

During the analysis phase of an SIR scheme, a weight must be assigned to each particle by comparison with the available observations. This requires computing an approximation to $p(\theta|\psi)$, the likelihood of recording an observation θ , given the modeled state ψ .

With SIR-HUXt, we are investigating the usefulness of assimilating the time-elongation profile of a CME flank which could be observed by HI-like instruments. Therefore, in this context, we must compute the likelihood of an observed CME flank elongation value for a specific modeled flank elongation.

Computation of the time-elongation profiles of the Cone CMEs is described in Section 2.1, whilst generation of the pseudo-observations, where Gaussian noise is added to the time-elongation profiles, is described in Section 2.3.

To compare the simulated and observed flank elongation, we use an assumed Gaussian likelihood profile for $p(\theta|\psi)$, which is centered on the simulated flank elongation with a spread of 0.15° . As far as we are aware, there is

no good a priori or empirical knowledge of how this likelihood profile should be structured. However, we believe that a Gaussian is a reasonable approximation, which can be refined in future.

Then, after computing the observation likelihood l_i for each ensemble member, its weight w_i is computed as the normalized likelihood over the N ensemble members,

$$w_i = \frac{l_i}{\sum_{i=1}^N l_i}. \quad (2)$$

These weights are then used in the resampling procedure.

2.2.3. Resampling

Kernel density estimation is used to compute the resampling of the ensemble members. Each Cone CMEs state is represented by its speed, width, and longitude. To resample the ensemble, we must draw samples from the kernel density estimate of the joint-distribution of these 3 parameters. However, scale separation of the parameters, particularly the CME speed from the width and longitude, means it is necessary to rescale the parameters before computing the kernel density estimate. Therefore, we compute the z-scores of each parameter before computing the kernel density estimate of the joint distribution and resampling, using a Gaussian kernel with a bandwidth of 0.2. This bandwidth value was arrived at experimentally. If the bandwidth is too large, the resampled ensemble will not be drawn toward the observations and so the assimilation achieves little to nothing. If it is too small, the resampled ensemble can be pulled too aggressively toward a highly weighted particle, making filter degeneracy more probable.

Figure 1 shows how the resampling procedure works in practice. The top row shows the Cone CME parameters of the initial ensemble (the prior) as red dots, for each pairing of the Cone CME speed, width, and longitude. The kernel density estimate of the distribution of these points is also contoured. These distributions are relatively uniform, given the sample size of 50, as would be expected from the uniform perturbation functions that generate the prior. In the bottom row the red dots show the same prior Cone CME parameters, but with a size proportional to their weight determined in the first analysis step of an SIR computation. Here, the contours instead show the distribution of the weighted prior Cone CME parameters. The black squares show the resampled Cone CME parameters that form the new ensemble that will be advanced to the next analysis step. It is clear that the new ensemble is closer to the prior Cone CME parameters that had larger weights, but is not overly concentrated around particles with the largest weights.

2.3. Observing System Simulation Experiment (OSSE)

An Observing System Simulation Experiment (OSSE) is a method with which we can assess the potential benefits of integrating a DA scheme into a physical model of a system (Zeng et al., 2020). OSSEs are controlled experiments using simulations of synthetic scenarios that allow us to explore the usefulness of different observation networks and/or DA schemes (Cucurull & Casey, 2021).

These experiments begin by using a model to simulate a ‘ground truth.’ Observations of this ground truth are generated by combining a forward model that emulates the observations from the ground truth with realistic observational noise. Then, these emulated observations are assimilated into the DA scheme, where the physical model is initialized with perturbed initial and/or boundary conditions relative to the ‘ground truth’ simulation. Through this process we can assess the ability of a DA scheme to recover the ‘ground truth.’

Here we use OSSEs configured as a ‘twin experiment.’ where we perform the same experiment with both HUXt and SIR-HUXt, to assess the performance of the SIR scheme relative to an ensemble of HUXt simulations without DA. This is the same general method as that employed by Lang et al. (2017), who investigated the use the Local Ensemble Transform Kalman Filter with the WSA-ENLIL solar wind model, and in Lang and Owens (2019), in the development of BRaVDA with the HUX solar wind model.

Figure 2 presents a flow diagram of the configuration of the OSSE experiments. To collect statistics on the performance of the SIR-HUXt scheme for a particular combination of CME scenario and Observer, the steps bounded in red are repeated 100 times, with different random realizations of: the guess at the CME initial conditions; the generation of the initial ensemble; and the observed time-elongation profile of the CME flank. The following subsections describe the individual steps in this flow chart.

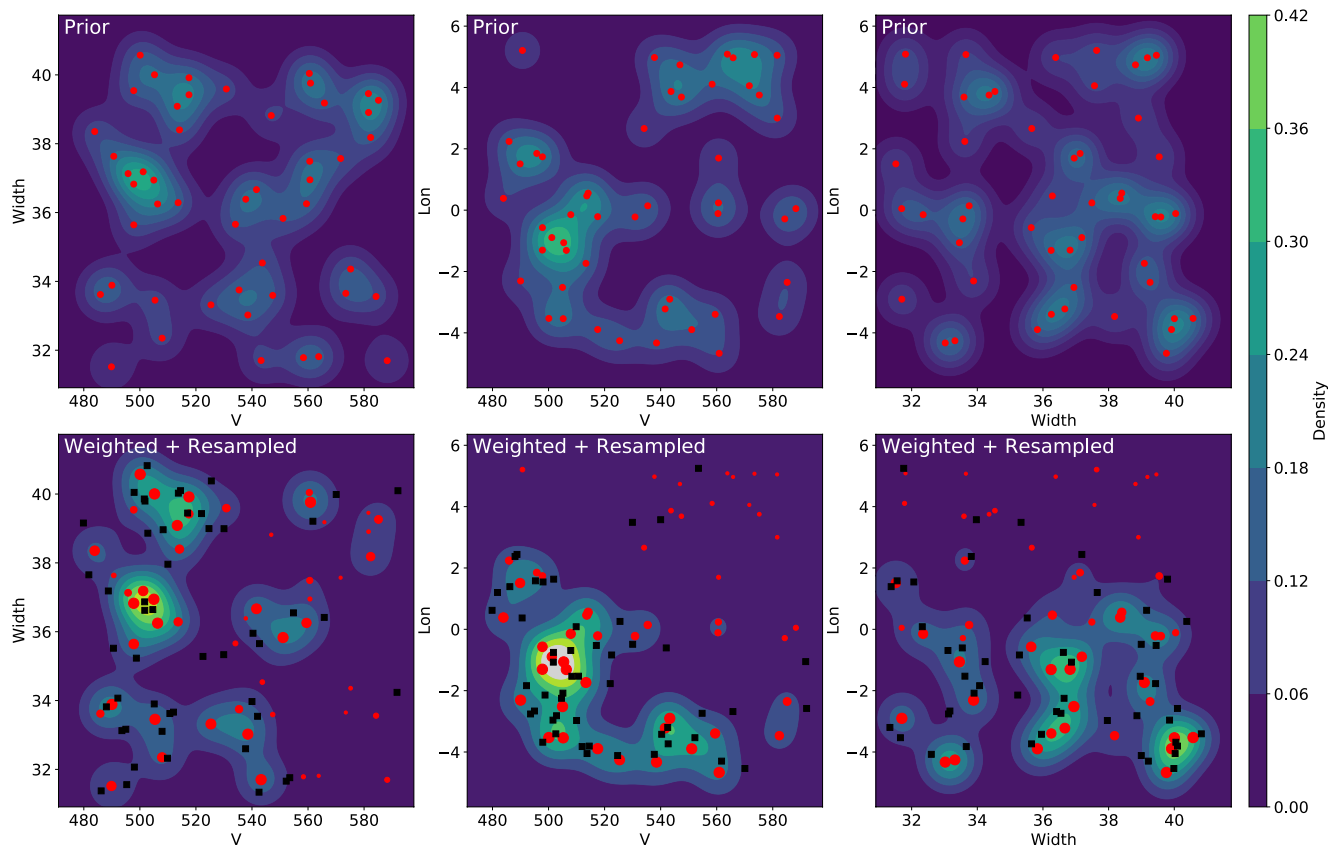


Figure 1. The top row shows the Cone Coronal Mass Ejection (CME) parameters of the initial ensemble (prior) in red dots, with contours of the kernel density estimates of their distribution. The bottom row shows the Cone CME parameters of the initial ensemble (prior) in red dots, with a size proportional to their weight determined in the sequential importance resampling analysis step. The contours show the kernel density estimates of these weighted distributions, whilst the black dots show the resampled Cone CME parameters that form the ensemble that will be advanced to the next analysis step.

2.3.1. CME Scenario

We use a Cone CME scenario to develop the SIR-HUXt system. This scenario reflects the climatological average CME. To build this scenario, we analyzed the distribution of observed CME speeds and widths provided by KINCAT database in the HELCATS project. The KINCAT data are described in D. Barnes et al. (2020) and Pluta et al. (2019), and consist of graduated cylindrical shell (GCS) fits (Thernisien, 2011) of 122 CMEs observed in the STEREO COR2 coronagraphs (Howard et al., 2008). These GCS fits return estimates of the CME apex speed and the angular half-width. These data are presented as a scatter plot in Figure 3.

We compute the medians of the CME speed and (full) width to define the average CME scenario. The Cone CME is fully Earth-directed, having the same source longitude and latitude as Earth, and is initialized 1 hr after the model start time. These values are shown by the orange hexagon in Figure 3 and summarized in the left panel of Figure 4.

A uniform ambient solar wind is used with the Cone CME scenario, with the ambient solar wind speed at the inner boundary being set to 400 km/s. We choose to use a uniform ambient solar wind in these experiments to reduce the complexity of the system whilst we develop SIR-HUXt. Future experiments will explore the impact of both different CME scenarios and structured solar wind on the performance of the SIR-HUXt. Figure 4 presents a snapshot from a HUXt simulation of the Cone CME scenario which. Provides the ground truth simulation against which the performance of the SIR scheme will be assessed.

2.3.2. HUXt Ground Truth Run

A HUXt simulation is produced using the unperturbed CME scenario, including calculations of the CME transit time to Earth and arrival speed, computed using the standard HUXt tools (M. Owens & Barnard, 2022).

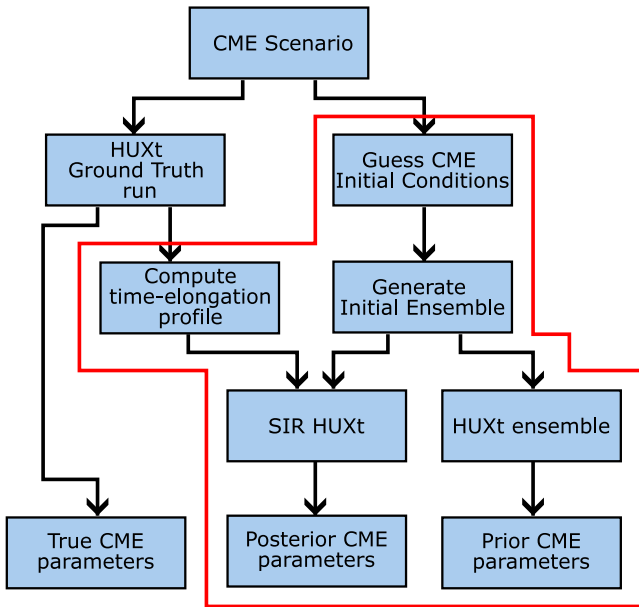


Figure 2. A flow chart showing the configuration of the Observing System Simulation Experiment experiments. The stages encapsulated in the red box are repeated for 100 iterations, with each iteration using a different random realization of the guessed Coronal Mass Ejection initial conditions, the computed time-elongation profile, and the initial ensemble.




Figure 3. The black dots mark the Coronal Mass Ejection (CME) speeds and widths from the HELCATS CME classifications determined from stereoscopic fits to coronagraph observations. The orange hexagon marks the speed and width parameters of the Cone CME scenario used in this study, which correspond to the median of the speed and width distributions, which are marked with red-dashed and teal-dotted lines.

2.3.3. Compute Time-Elongation Profile

An observer tracks the time-elongation profile of the Cone CME flank in the HUXt ground truth run, as described in Section 2.1. The flank is tracked over the elongation range spanning 4° – 35° , with observations recorded every 174 min (corresponding to 20 steps of HUXt's native time-step). To these time-elongation profiles, Gaussian noise is added to the elongations with a standard deviation of 0.1° . We consider this a reasonable lower limit on the elongation uncertainty, as analysis of time-elongation profiles extracted from STEREO Heliospheric Imager data suggest that elongation uncertainties of $\approx 0.5^{\circ}$ are typical (Barnard et al., 2015, 2017; Möstl et al., 2011; Williams et al., 2009).

In each SIR-HUXt OSSE, only observations from one observer are assimilated. However, to investigate the impact of observer longitude, we run the experiments with 8 observers at longitudes spanning -90° to -20° in steps of 10° , all situated in the same latitudinal plane as Earth. Figure 4 panel c shows the locations of the 8 observers that track the time-elongation profiles of the Cone CMEs. Panel b also shows the field-of-view of the observer (red shaded region) situated at -60° , corresponding to the L5 location.

2.3.4. Guess CME Initial Conditions

In both research and forecasting simulations of real world CME's, we do not have perfect knowledge of a CME's initial conditions. These parameters must be estimated from observations and/or empirical relations. To emulate this process here, in each OSSE, we make a guess at the CME initial conditions by applying a perturbation to the ground truth Cone CME parameters of the CME scenario. The perturbations are calculated using the same procedure as is used to generate the initial ensemble of Cone CME parameters, as described in Section 2.2.1. To summarize, these are perturbations to the Cone CME speed, width, and longitude, derived from uniform distributions that approximate the uncertainties on the estimated CME parameters.

2.3.5. Generate Initial Ensemble

The initial ensemble of Cone CME parameters, which is used in both the SIR-HUXt simulations and the HUXt ensemble, is generated according to the procedure described in 2.2.1. In the generation of the initial ensemble, the guess of the CMEs initial conditions is used as the best estimate to which the perturbations are applied.

2.3.6. SIR-HUXt

SIR-HUXt takes the initial ensemble and the observed time-elongation profile and performs eight iterations of the SIR analysis. These eight analysis steps are the maximum that can be performed consistently across all experiments, within the observers constraints of the field-of-view extending to only 35° elongation, and recording observations every 174 min. At each analysis step, the Cone CME parameters of each ensemble member are recorded, as are the CME transit time to Earth and arrival speed, which is computed using the standard HUXt tools.

2.3.7. HUXt Ensemble

The HUXt ensemble run proceeds by simply generating a HUXt simulation for each member of the initial ensemble. Similarly, each of the Cone CME parameters are recorded, as are the CME transit time to Earth and arrival speed, computed using the standard HUXt tools.




Figure 4. (left) A snapshot of a HUXt simulation of the median Coronal Mass Ejection (CME) scenario approximately 1 day after model initialization. The background solar wind speed is uniform, being 400 km/s at the model inner boundary. Earth is marked by the white dot, whilst the CME boundary is marked by the orange line. The 8 pseudo-observers used in this study are marked by the black squares, with the exception of the L5 observer, marked by the red square. The field-of-view of the L5 observer covers the area bounded between the red dotted lines, while the elongation of the tracked flank of the Cone CME is shown by the red dashed line. (right) The black line shows the time-elongation profile of the Cone CME flank for the L5 observer. The red squares mark the time-elongation coordinates with added Gaussian noise to emulate the uncertainty introduced by extracting the coordinates from heliospheric imager data, as described in Section 2.3.3.

2.3.8. The True, Prior, and Posterior CME Parameters

From the above simulations we have knowledge of the true CME parameters, including transit time to Earth and arrival speed, as well as the prior distributions of these parameters, returned by the HUXt ensemble, and the posterior distributions, returned by SIR-HUXt. These data are then used in the statistical assessment of the performance of SIR-HUXt relative to a simple ensemble of HUXt runs.

3. Results

3.1. An L5 Observer of the Median CME Scenario

This Observer-CME scenario combination is a highly relevant scenario for future space weather forecasting, for two reasons. First, the median CME scenario is a robust estimate of the average CME properties. Second, as ESA's Vigil mission will provide heliospheric imaging data from L5 for use in operational space weather forecasts, the L5 Observer approximates the time-elongation profiles that might be obtained operationally by Vigil's HI.

3.1.1. Example of One Realization of SIR-HUXt Analysis

Figure 5 presents an example of a SIR-HUXt analysis, showing how the ensemble evolves as a function of the number of analysis steps. These data are from a single SIR-HUXt analysis from the 100 realizations in the OSSE experiment. In this instance, the initial guess of the CME longitude, width and speed was -4° , 37° , and 477 km s^{-1} , around which the initial ensemble (analysis step 0) was formed. For the longitude, speed, and transit time, the distributions evolve significantly over the analysis steps, both moving toward and reducing in spread around the true value. In this example the width distribution is less strongly impacted by the SIR analysis, drifting slightly whilst maintaining a similar spread.

Considering panel d, the uncertainty in the CME transit time (or correspondingly, arrival time), is significantly reduced from 5.3 hr in the initial ensemble to 0.9 hr after the SIR-HUXt analysis. These transit time errors are




Figure 5. These panels show the evolution of the SIR-HUXt ensemble as a function of the number of analysis steps taken. Panels (a, b, c, and d) show the evolution of the longitude, width, speed, and transit time distributions, respectively. In each panel, the gray lines mark each of the 50 ensemble members, while the blue, orange, and green lines mark the 10th percentile, median, and 90th percentiles. The red dashed line marks the true value of each parameter.

smaller than the average observed transit time errors; this is primarily due to idealized scenario used in our experiment, which does not yet include ambient solar wind structure.

We stress that this is only one example, and that alternative behaviors are observed. Additionally, it is not “wrong” or a failure of the SIR scheme that the width distribution does not change much during the analysis. Depending on the initial estimate of the CME parameters, and the uncertainty on the observations, it is possible that the distribution of any particular parameter need not evolve significantly. Additionally, in this observing geometry, it is reasonable to expect that the observed CME flank will be less sensitive to the CME width than the longitude and speed, particularly in the absence of structured solar wind. For different viewing geometries, for example, halo CMEs, we would expect the CME width to more significantly impact the time-elongation profiles of the CME flank and hence the SIR HUXt results.

3.1.2. Aggregated SIR-HUXt OSSE Results

Figure 6 compares the prior and posterior distributions of the CME parameters, aggregated over all of the OSSE experiments. These are presented as three 2-D histograms, showing the joint distributions of the CME speed and width, speed and longitude, and width and longitude. The top row shows the prior distributions, whilst the bottom row shows the posterior distributions. The red dashed lines mark the true parameter values. As each SIR-HUXt run uses a 50 member ensemble, and there are 100 realizations in the OSSE, there are 5,000 samples in each distribution.

The prior distributions are relatively uniform, as expected from the perturbation function used to generate the initial ensembles. The posterior distributions have significantly different structure to the priors. Panels d and e show that the speed distribution has been strongly constrained around the true value, with the standard deviation reducing from 40 km s^{-1} to 11 km s^{-1} . The standard deviation in width distribution is approximately the same for the prior and posterior, being 4.1° and 4.3° . The standard deviation of the posterior longitude distribution is reduced relative to the prior, decreasing from 4.1° to 3.0° . Panel e also shows the emergence of a correlation between the posterior distributions of speed and longitude. This is not surprising, as the time-elongation profiles




Figure 6. These panels show 2-D histograms of the prior and posterior distributions of the Coronal Mass Ejection parameters for an L5 Observer. The prior distributions are shown along the top row, with the posterior distributions on the bottom row. The red dashed lines mark the true parameter values.

of CME flanks have a well known degeneracy relating to the speed, width and longitude. Similar time-elongation profiles can be generated by CMEs traveling at different angles relative to the plane-of-sky with different speeds. This degeneracy can lead to such correlations when trying to find a combination of CME parameters that best reflects a time-elongation measurement. Consequently, we expect this to be a “feature” of assimilating CME time-elongation profiles from one observer. We note that it is possible that assimilating HI data from more than one observer, or, assimilating the HI image intensities rather than only a time-elongation profile, might break these degeneracies, and these objectives should be a priority for future investigation. Nonetheless, it is clear that the posterior CME states are typically closer to the true CME state, even if this is dominated by the evolution of the CME speed distribution.

3.1.3. Ensemble Mean SIR-HUXt OSSE Results

It is also instructive to compare the means of the prior and posterior distributions for each realization of the OSSE experiment. Figure 7 shows these data, using the same format as Figure 6. Each histogram contains 100 samples from the 100 OSSE experiments, with each 50-member ensemble reduced to its mean value. We observe the same behavior in these distributions as was observed for the aggregated SIR-HUXt analyses in Figure 6. The CME speeds are strongly constrained around the true value, with the standard deviation reducing from 28 km s^{-1} to 8 km s^{-1} . There are only small changes between the prior and posterior distributions of the CME width and longitude. The spread of the distribution of CME widths increases slightly with a prior standard deviation of 2.9° and posterior standard deviation of 3.4° . Conversely, the spread of the CME longitudes decreases slightly, with a prior standard deviation of 2.9° and posterior standard deviation of 2.2° .

3.1.4. Assessment of Ensemble Representivity With Rank-Histograms

It is also important to assess the representivity of the SIR-HUXt ensembles. If the SIR-HUXt ensembles were perfectly calibrated, then each ensemble member and the truth state would be independent samples from the




Figure 7. These panels show 2-D histograms of the mean Coronal Mass Ejection parameters for each SIR-HUXt realization of the Observing System Simulation Experiment experiment. The format is the same as Figure 6.

same underlying probability distribution. A rank histogram is a graphical means of assessing this (Talagrand et al., 1997). To construct the rank histogram we rank the true system state in each SIR-HUXt realization, and plot a histogram of these data. If the truth state and ensemble members are independent samples from the same probability distribution, then the rank histogram would be uniform, to within the limits of sampling variability. However, deviations from uniformity can diagnose miscalibrations in the ensemble. For example, if the ensemble is over or under dispersed, the rank-histogram takes a U or inverted-U shape, or if the ensemble is biased the rank histogram can be asymmetric (Wilks, 2019).

Figure 8 presents the rank histograms for the prior and posterior distributions of CME speed, width, and longitude. As the prior distributions are generated by uniform perturbations to the CME scenario, it follows that the rank-histogram of the prior distributions are also uniform, to within the limits of sampling variability. It is clear that the posterior distributions show a similar level of uniformity, which is one indicator that the SIR-HUXt ensembles are reasonably well calibrated.

3.1.5. SIR-HUXt Impact on CME Transit Time and Arrival Speed Distributions

Finally, we consider the impact of SIR-HUXt on the distributions of CME transit time and arrival speed at Earth. Figure 9 presents these data. Panel a and b show histograms of the prior and posterior distributions of CME transit time and arrival speed at Earth. Panel c shows a scatter plot of the CME transit time versus the CME arrival speed for both the prior and posterior distributions. For this scenario, the true CME transit time and arrival speed were 70.8 hr, and 498 km s⁻¹. Considering panel a, the prior distribution has a larger spread around the true transit time than the posterior distribution. The standard deviation of the prior and posterior distributions are 2.5 and 0.8 hr, respectively. Therefore the SIR-HUXt analysis results in a 69% reduction in the CME transit time standard deviation. Panel b shows similar results for prior and posterior distributions of the CME arrival speed; the posterior distribution is less spread around the true value than the prior, with standard deviations of 11 km s⁻¹ and 4 km s⁻¹, respectively, a 63% reduction in the arrival speed standard deviation. Panel c shows that, as expected




Figure 8. Rank histograms of the prior and posterior distributions of the Coronal Mass Ejection speed, width, and longitude.

for this scenario without ambient solar wind structure, there is a clear anti-correlation between CME transit time and arrival speed at Earth. This correlation is present for both the joint-posterior and joint-prior distributions. However, consistent with panels a and b, the joint-posterior distribution is closer to the true CME transit time and arrival speed for this scenario. In this way, we consider this evidence that SIR-HUXt has significant potential for improving CME transit time and arrival speed simulations over simple HUXt ensembles.

3.2. SIR-HUXt Performance With Observer Longitude

We now consider how the posterior distributions of the Cone CME parameters and the transit time and arrival speed at Earth vary as a function of observer longitude. These data are presented in Figure 10, where the prior and posterior distributions of each parameter are summarized by their lower decile, median, and upper decile. The true parameter values are shown by the red dashed line.

The prior distributions show no variation with observer longitude, because the same set of randomly generated initial ensembles are used for the SIR-HUXt OSSE experiments at each longitude, so as to enable a fair comparison between longitudes. For each parameter, the median of the prior and posterior distributions are very similar and close to the true parameter values, indicating no significant bias between the prior and posterior distributions with the true parameter values.

There are, however, systematic changes in the spread of the SIR-HUXt posterior parameters as a function of observer longitude, where we define spread to be the difference between the lower and upper deciles. For the




Figure 9. Panels (a and b) show histograms of the prior and posterior distributions of the Coronal Mass Ejection (CME) transit time and arrival speed, respectively. Panel (c) presents a scatter plot of the CME transit time and arrival speed. The true CME transit time and arrival speed are shown by the red dashed lines.




Figure 10. These panels show the evolution of prior and posterior distributions of the Cone Coronal Mass Ejection speed, longitude, width, transit time, and arrival speed, as a function of observer longitude. These distributions are summarized by their lower decile (triangles), median (circles) and upper decile (crosses), with the priors and posteriors colored blue and orange, respectively. The true parameter values are shown by the red dashed line.

initial Cone CME speed, the posterior spread is less than the prior spread at all observing longitudes, and it also shows a local minimum at 290° longitude. This indicates that the SIR-HUXt posteriors provide a tighter constraint on the CME speed from all observer longitudes, but the tightest constrain comes from assimilating time-elongation profiles from observers close to the L5 region. We think this behavior is driven by the fact that for an observer in this region, with a fully Earth directed CME, the time-elongation profile of the flank corresponds closely to the CME apex (see Figure 4 panel b), which minimizes the degeneracy between the CME speed, longitude and width. However, this is not the case for the CME longitude, where the spread of the distribution continues to increase as the observer moves from 340° to 290° . This suggests that SIR-HUXt is better able to constrain the CME source longitude from Observers nearer Earth. There seem to be no significant differences

between the prior and posterior distributions of the CME width, suggesting that for this particular scenario SIR-HUXt does not have a significant impact on the CME width estimation. Further experiments are required to determine whether this behavior is general, or is specific to this particular CME scenario. Both the CME transit time and arrival speed show behavior that mirrors that of the CME speed, with the spread being less than the prior for all observing longitudes, and showing a local minimum at 290° . This is unsurprising, given that for this CME scenario, with uniform ambient solar wind, we expect the transit time and arrival speed to be primarily determined by the Cone CME initial speed.

4. Conclusions

In this work we have presented the development of SIR-HUXt, a particle filter DA scheme for constraining the HUXt solar wind model. SIR-HUXt assimilates time-elongation profiles of a CMEs flank, which is a data product routinely derived from HI data, such as that returned by STEREO-HI, and Parker Solar Probe's WISPR.

The motivation for pursuing the development of SIR-HUXt is that, at present, there is significant uncertainty in the initial and boundary conditions of the solar wind numerical models that are used for both scientific and forecasting purposes. These uncertainties limit both the scientific inferences and forecast skill we can extract from solar wind simulations. DA techniques present a pathway to reduce these uncertainties, providing a framework for combining simulations with observations to return an optimal estimate of the state of the system. HUXt is well suited to the development of DA schemes due to its simplicity and low computational expense. This latter point is particularly important for the development of this Sequential Importance Resample particle filter, which requires a large ensemble of simulations to be run (≈ 500 per SIR-HUXt analysis); this would be very computationally expensive for full 3D MHD solar wind models.

In its current form, SIR-HUXt adjusts the Cone CME parameters in response to assimilating the time-elongation profile of an observed CMEs flank. We used OSSEs to provide an initial proof-of-concept test of the SIR-HUXt algorithm. These experiments test the ability of SIR-HUXt to recover a known synthetic truth state, by assimilating pseudo-observations of the known truth state. In these experiments, our truth state was a simple scenario of an Earth directed CME with the median speed and width of observed CMEs, propagating through a uniform ambient solar wind. These OSSEs show that SIR-HUXt is effective at constraining the CME state, primarily by adjusting the CME speed. These experiments suggest SIR-HUXt was less effective at constraining the CME longitude and width. This is likely due to the fact the separating the effects of CME longitude and width is an inherently uncertain process in single-point HI observations, as time elongation profiles of the CME flank can be essentially degenerate for a range of longitude and width values. We expect this issue would be improved by the assimilation of HI data from multiple longitudinally separated observatories.

Nonetheless, by improving the constraint of the initial CME state, SIR-HUXt also returns improved estimates of the CME transit time to Earth, and the CMEs arrival speed, which are critical parameters for space weather predictions. The reliability of the SIR-HUXt ensembles was assessed through rank-histogram plots, through which we conclude that the SIR-HUXt ensembles are reasonably well calibrated, with no clear indications of under or over dispersion, or bias.

The OSSEs also revealed that the location of the observer relative to the CME has a significant impact on the ability of SIR-HUXt to constrain the CMEs parameters. Observers near the L5 location provided the best constraints on the CME speed, transit time and arrival speed. The SIR-HUXt constraints on the CME longitude grew weaker as separation between the observers longitude and CME apex longitude increased. This is a potentially significant result relating to likely performance of schemes like SIR-HUXt for space weather forecasting with the operational HI data that will be returned by ESA's Vigil mission.

We note that it is now well established that both the CME initial conditions and the ambient solar wind structure both play important roles in determining the simulated CME evolution and, critically, the forecast arrival time at Earth. Indeed, Riley and Ben-Nun (2021) investigated the sources of uncertainty in CME arrival time predictions and concluded that both the ambient solar wind structure and CME parameters introduce similar magnitudes of uncertainty. Our work has so far only considered uncertainty in the CME parameters, and does not yet consider uncertainty in the ambient solar wind structure, although we have plans to investigate this issue in future work.

Following this study, our next two objectives are to test SIR-HUXt with OSSEs using a wider range of CME scenarios, and to test SIR-HUXt with actual time-elongation profiles extracted from the STEREO-HI data.

Data Availability Statement

The software and data required to generate all of the simulation results used in this study are archived with Zenodo at <https://doi.org/10.5281/zenodo.7962973> and are also available at: https://github.com/University-of-Reading-Space-Science/SIR_HUXt.

Acknowledgments

This research made use of Astropy (<http://www.astropy.org>), a community-developed core Python package for Astronomy (Price-Whelan et al., 2018; Robitaille et al., 2013). This research used version 4.0.0 (Mumford et al., 2022) of the SunPy open source software package (W. T. Barnes et al., 2020). Figures for this article were made with version 3.3.4 of Matplotlib (Caswell et al., 2021; Hunter, 2007). This work was part-funded by Science and Technology Facilities Council (STFC) Grant ST/R000921/1 and ST/V000497/1, and NERC Grant NE/S010033/1.

References

- Ades, M., & van Leeuwen, P. J. (2013). An exploration of the equivalent weights particle filter. *Quarterly Journal of the Royal Meteorological Society*, 139(672), 820–840. <https://doi.org/10.1002/qj.1995>
- Arge, C. N., & Pizzo, V. J. (2000). Improvement in the prediction of solar wind conditions using near-real time solar magnetic field updates. *Journal of Geophysical Research*, 105(A5), 10465–10479. <https://doi.org/10.1029/1999JA000262>
- Barnard, L., de Koning, C. A., Scott, C. J., Owens, M. J., Wilkinson, J., & Davies, J. A. (2017). Testing the current paradigm for space weather prediction with heliospheric imagers. *Space Weather*, 15(6), 782–803. <https://doi.org/10.1002/2017SW001609>
- Barnard, L., & Owens, M. (2022). HUXt—An open source, computationally efficient reduced-physics solar wind model, written in Python. arXiv. <https://doi.org/10.48550/arXiv.2210.00455>
- Barnard, L., Owens, M. J., Scott, C. J., & de Koning, C. A. (2020). Ensemble CME modeling constrained by heliospheric imager observations. *AGU Advances*, 1(3), e2020AV000214. <https://doi.org/10.1029/2020AV000214>
- Barnard, L., Owens, M. J., Scott, C. J., Lockwood, M., de Koning, C. A., Amerstorfer, T., et al. (2021). Quantifying the uncertainty in CME kinematics derived from geometric modeling of heliospheric imager data. *Space Weather*, 20(1), e2021SW002841. <https://doi.org/10.1029/2021SW002841>
- Barnard, L., Scott, C. J., Owens, M. J., Lockwood, M., Crothers, S. R., Davies, J. A., & Harrison, R. A. (2015). Differences between the CME fronts tracked by an expert, an automated algorithm, and the Solar Stormwatch project. *Space Weather*, 13(10), 709–725. <https://doi.org/10.1002/2015SW001280>
- Barnes, D., Davies, J. A., Harrison, R. A., Byrne, J. P., Perry, C. H., Bothmer, V., et al. (2020). CMEs in the heliosphere: III. A statistical analysis of the kinematic properties derived from stereoscopic geometrical modelling techniques applied to CMEs detected in the heliosphere from 2008 to 2014 by STEREO/HI-1. *Solar Physics*, 295(11), 150. <https://doi.org/10.1007/s11207-020-01717-w>
- Barnes, W. T., Bobra, M. G., Christe, S. D., Freij, N., Hayes, L. A., Ireland, J., et al. (2020). The SunPy project: Open source development and status of the version 1.0 core package. *The Astrophysical Journal*, 890(1), 68. <https://doi.org/10.3847/1538-4357/ab4f7a>
- Browne, P. A., & van Leeuwen, P. J. (2015). Twin experiments with the equivalent weights particle filter and HadCM3. *Quarterly Journal of the Royal Meteorological Society*, 141(693), 3399–3414. <https://doi.org/10.1002/qj.2621>
- Bunting, K. A., & Morgan, H. (2022). An inner boundary condition for solar wind models based on coronal density. *Journal of Space Weather and Space Climate*, 12, 30. <https://doi.org/10.1051/swsc/20220206>
- Burgers, G., van Leeuwen, P. J., & Evensen, G. (1998). Analysis scheme in the ensemble Kalman filter. *Monthly Weather Review*, 126(6), 1719–1724. [https://doi.org/10.1175/1520-0493\(1998\)126<1719:ASITEK>2.0.CO;2](https://doi.org/10.1175/1520-0493(1998)126<1719:ASITEK>2.0.CO;2)
- Caswell, T. A., Droettboom, M., Lee, A., de Andrade, E. S., Hunter, J., Firing, E., et al. (2021). Matplotlib/matplotlib: REL: V3.3.4. Zenodo. <https://doi.org/10.5281/zenodo.4475376>
- Chi, Y., Scott, C., Shen, C., Barnard, L., Owens, M., Xu, M., et al. (2021). Modeling the observed distortion of multiple (ghost) CME fronts in STEREO heliospheric imagers. *The Astrophysical Journal Letters*, 917(2), L16. <https://doi.org/10.3847/2041-8213/ac1203>
- Chorin, A. J., & Tu, X. (2009). Implicit sampling for particle filters. *Proceedings of the National Academy of Sciences of the United States of America*, 106(41), 17249–17254. <https://doi.org/10.1073/pnas.0909196106>
- Cucurull, L., & Casey, S. P. F. (2021). Improved impacts in observing system simulation experiments of radio occultation observations as a result of model and data assimilation changes. *Monthly Weather Review*, 149(1), 207–220. <https://doi.org/10.1175/MWR-D-20-0174.1>
- Davies, J. A., Harrison, R. A., Rouillard, A. P., Sheeley, N. R., Perry, C. H., Bewsher, D., et al. (2009). A synoptic view of solar transient evolution in the inner heliosphere using the heliospheric imagers on STEREO. *Geophysical Research Letters*, 36(2), L02102. <https://doi.org/10.1029/2008GL036182>
- Eyles, C. J., Harrison, R. A., Davis, C. J., Waltham, N. R., Shaughnessy, B. M., Mapson-Menard, H., et al. (2008). The heliospheric imagers onboard the STEREO mission. *Solar Physics*, 254(2), 387–445. <https://doi.org/10.1007/s11207-008-9299-0>
- Fallow, R. A., Iwai, K., Jackson, B. V., Zhang, P., Bisi, M. M., & Zucca, P. (2022). Application of novel interplanetary scintillation visualisations using LOFAR: A case study of merged CMEs from September 2017. *Advances in Space Research*. <https://doi.org/10.1016/j.asr.2022.08.076>
- Fearnhead, P., & Künsch, H. R. (2017). Particle filters and data assimilation. arXiv. <https://doi.org/10.48550/arXiv.1709.04196>
- Howard, R. A., Moses, J. D., Vourlidas, A., Newmark, J. S., Socker, D. G., Plunkett, S. P., et al. (2008). Sun Earth connection coronal and heliospheric investigation (SECCHI). *Space Science Reviews*, 136(1–4), 67–115. <https://doi.org/10.1007/s11214-008-9341-4>
- Hunter, J. D. (2007). Matplotlib: A 2D graphics environment. *Computing in Science & Engineering*, 9(3), 90–95. <https://doi.org/10.1109/MCSE.2007.55>
- Iwai, K., Shiota, D., Tokumaru, M., Fujiki, K., Den, M., & Kubo, Y. (2021). Validation of coronal mass ejection arrival-time forecasts by magnetohydrodynamic simulations based on interplanetary scintillation observations. *Earth Planets and Space*, 73(1), 9. <https://doi.org/10.1186/s40623-020-01345-5>
- Lang, M., Browne, P., van Leeuwen, P. J., & Owens, M. J. (2017). Data assimilation in the solar wind: Challenges and first results. *Space Weather*, 15(11), 1490–1510. <https://doi.org/10.1002/2017SW001681>
- Lang, M., & Owens, M. J. (2019). A variational approach to data assimilation in the solar wind. *Space Weather*, 17(1), 59–83. <https://doi.org/10.1029/2018SW001857>
- Lang, M., Witherington, J., Turner, H., Owens, M. J., & Riley, P. (2021). Improving solar wind forecasting using data assimilation. *Space Weather*, 19(7), e2020SW002698. <https://doi.org/10.1029/2020SW002698>
- Le Dimet, F.-X., & Talagrand, O. (1986). Variational algorithms for analysis and assimilation of meteorological observations: Theoretical aspects. *Tellus A: Dynamic Meteorology and Oceanography*, 38(2), 97–110. <https://doi.org/10.3402/tellusa.v38i2.11706>
- Möstl, C., Rollett, T., Lugaz, N., Farrugia, C. J., Davies, J., Temmer, M., et al. (2011). Arrival time calculation for interplanetary coronal mass ejections with circular fronts and application to stereo observations of the 2009 February 13 eruption. *The Astrophysical Journal*, 741(1), 34. <https://doi.org/10.1088/0004-637X/741/1/34>
- Mumford, S. J., Freij, N., Stansby, D., Christe, S., Ireland, J., Mayer, F., et al. (2022). SunPy. Zenodo. <https://doi.org/10.5281/zenodo.6524764>

- Owens, M., & Barnard, L. (2022). University-of-Reading-Space-Science/HUXt: HUXt 4.0. *Zenodo*. <https://doi.org/10.5281/zenodo.6794462>
- Owens, M. J., Lang, M., Barnard, L., Riley, P., Ben-Nun, M., Scott, C. J., et al. (2020). A computationally efficient, time-dependent model of the solar wind for use as a surrogate to three-dimensional numerical magnetohydrodynamic simulations. *Solar Physics*, 295(3), 43. <https://doi.org/10.1007/s11207-020-01605-3>
- Pant, V., Willems, S., Rodriguez, L., Mierla, M., Banerjee, D., & Davies, J. A. (2016). Automated detection of coronal mass ejections in stereo heliospheric imager data. *The Astrophysical Journal*, 833(1), 1–15. <https://doi.org/10.3847/1538-4357/833/1/80>
- Pizzo, V. J., de Koning, C. A., Cash, M., Millward, G., Biesecker, D. A., Puga, L., et al. (2015). Theoretical basis for operational ensemble forecasting of coronal mass ejections. *Space Weather*, 13(10), 676–697. <https://doi.org/10.1002/2015SW001221>
- Pluta, A., Mrotzek, N., Vourlidas, A., Bothmer, V., & Savani, N. (2019). Combined geometrical modelling and white-light mass determination of coronal mass ejections. *Astronomy & Astrophysics*, 623, A139. <https://doi.org/10.1051/0004-6361/201833829>
- Pothast, R., Walter, A., & Rhodin, A. (2019). A localized adaptive particle filter within an operational NWP framework. *Monthly Weather Review*, 147(1), 345–362. <https://doi.org/10.1175/MWR-D-18-0028.1>
- Price-Whelan, A. M., Sipőcz, B. M., Günther, H. M., Lim, P. L., Crawford, S. M., Conseil, S., et al. (2018). The Astropy project: Building an open-science project and status of the v2.0 core package. *The Astronomical Journal*, 156(3), 123. <https://doi.org/10.3847/1538-3881/aabc4f>
- Riley, P., & Ben-Nun, M. (2021). On the sources and sizes of uncertainty in predicting the arrival time of interplanetary coronal mass ejections using global MHD models. *Space Weather*, 19(6), e2021SW002775. <https://doi.org/10.1029/2021SW002775>
- Riley, P., Linker, J. A., & Mikić, Z. (2001). An empirically-driven global MHD model of the solar corona and inner heliosphere. *Journal of Geophysical Research*, 106(A8), 15889–15901. <https://doi.org/10.1029/2000JA000121>
- Riley, P., & Lionello, R. (2011). Mapping solar wind streams from the Sun to 1 AU: A comparison of techniques. *Solar Physics*, 270(2), 575–592. <https://doi.org/10.1007/s11207-011-9766-x>
- Riley, P., Mays, M. L., Andries, J., Amerstorfer, T., Biesecker, D., Delouille, V., et al. (2018). Forecasting the arrival time of coronal mass ejections: Analysis of the CCMC CME scoreboard. *Space Weather*, 16(9), 1245–1260. <https://doi.org/10.1029/2018SW001962>
- Robitaille, T. P., Tollerud, E. J., Greenfield, P., Droettboom, M., Bray, E., Aldcroft, T., et al. (2013). Astropy: A community Python package for astronomy. *Astronomy & Astrophysics*, 558, A33. <https://doi.org/10.1051/0004-6361/201322068>
- Talagrand, O., Vautard, R., & Strauss, B. (1997). *Evaluation of probabilistic prediction systems*. In (pp. 1–26). Shinfield Park, Reading: ECMWF.
- Thernisien, A. (2011). Implementation of the graduated cylindrical shell model for the three-dimensional reconstruction of coronal mass ejections. *The Astrophysical Journal - Supplement Series*, 194(2), 33. <https://doi.org/10.1088/0067-0049/194/2/33>
- Van Leeuwen, P. J. (2009). Particle filtering in geophysical systems. *Monthly Weather Review*, 137(12), 4089–4114. <https://doi.org/10.1175/2009MWR2835.1>
- Wilks, D. S. (2019). Indices of rank histogram flatness and their sampling properties. *Monthly Weather Review*, 147(2), 763–769. <https://doi.org/10.1175/MWR-D-18-0369.1>
- Williams, A. O., Davies, J. A., Milan, S. E., Rouillard, A. P., Davis, C. J., Perry, C. H., & Harrison, R. A. (2009). Deriving solar transient characteristics from single spacecraft STEREO/HI elongation variations: A theoretical assessment of the technique. *Annales Geophysicae*, 27(12), 4359–4368. <https://doi.org/10.5194/angeo-27-4359-2009>
- Yeates, A. R., Mackay, D. H., van Ballegoijen, A. A., & Constable, J. A. (2010). A nonpotential model for the Sun's open magnetic flux. *Journal of Geophysical Research*, 115(A9), A09112. <https://doi.org/10.1029/2010JA015611>
- Zeng, X., Atlas, R., Birk, R. J., Carr, F. H., Carrier, M. J., Cucurull, L., et al. (2020). Use of Observing system simulation experiments in the United States. *Bulletin of the American Meteorological Society*, 101(8), E1427–E1438. <https://doi.org/10.1175/BAMS-D-19-0155.1>